Modeling progress: causal models and the imperfective paradox

Elitzur Bar-Asher Siegal
Hebrew University of Jerusalem
ebas@mail.huji.ac.il

Prerna Nadathur
The Ohio State University
nadathur.3@osu.edu

June 20, 2026

Abstract

Durative telic *(accomplishment)* predicates are usually taken to denote sets of complete eventualities, each including realization of a lexically-specified culmination. This assumption is in tension with the observation that accomplishment predicates can be used to describe **non-culminated** events: most famously, in the scope of *progressive* aspects. Prominent accounts of this **imperfective paradox** (Dowty, 1979) tie the effect to the modal accessibility of culmination, intensionalizing the progressive so that it instantiates culminated events in certain alternatives to the evaluation world. We argue that this approach faces empirical challenges, combining data from progressives of *unlikely events* and locally *out of reach* tasks with insights from extensional treatments of the paradox to motivate a new analysis, on which telic progressives are sensitive to part-whole structure inherited from an **event type** associated with the underlying predicate. We propose that an event type corresponds to a **causal model** (Pearl, 2000) in which the *culmination condition* C_P of P occurs as a dependent variable. The model provides a set of *procedures* (causal pathways) for realizing C_P, each comprising a set of jointly sufficient causal conditions for C_P; it also establishes sets of conditions that preclude C_P. This approach allows the interpretation of telic progressives to depend not on modal projections of culmination, but instead on a partial match between reference time facts and sets defined by the event type: concretely, a reference situation verifies an *in-progress* P-eventuality just in case it satisfies a non-trivial subset of the conditions in some causal pathway for C_P (initiating a process for culmination), and fails to verify a sufficient set of conditions for non-culmination (rendering further development towards C_P possible). While our approach complexifies the denotations of telic predicates, it offers a simplified *partitive* view of grammatical aspects. Looking ahead, the account suggests a link between lexical aspectual properties and the structure of causal models which promises to shed light on (non)culmination phenomena beyond progressive contexts.

1 Introduction

The prevalent approach to verb meaning in contemporary lexical semantics is one of *predicate decomposition* (Hale and Keyser, 2002; Grimshaw, 2005; Levin and Rappaport Hovav, 2005, a.o.). On one prominent version of this approach, verb meaning is split into two components, one structural and the other idiosyncratic.[1] The first component characterizes *event type*, dividing verbs

[1] Although this paper will motivate an alternative to bipartite decomposition by examining a particular interaction between verb meaning and grammatical aspect—Dowty's (1979) **imperfective paradox**—we begin by presenting the bipartite approach because it represents an influential line of analysis from which our own proposal for verb meaning diverges in clear and specific ways. For a alternative perspective on decomposition (and one which may ultimately be more compatible with our approach), see, e.g., Ramchand (2008).

into semantic clusters or *classes* on the basis of the inherent internal structure of the events they describe, and feeding the grammatical properties of the verb. The idiosyncratic *root* component differentiates verbs within a structural class, imparting specific but grammatically inert semantic content. Within a bipartite decompositional architecture, roots are systematically linked to event structures: each root has an ontological classification drawn from a fixed spectrum of types (e.g., *state, result state, location, manner,* etc.) which determines where and how it can be integrated into a given event schema. The resulting representations are typically named by their roots, while the event scheme—expressed as a combination of primitive predicates such as ACT, CAUSE, BECOME (Dowty, 1979)—classifies the verb as either simple (comprising a single subevent) or complex (comprising two subevents; Levin and Rappaport Hovav 1999). Example (1) provides simple schema, corresponding respectively to *activities, states,* and *achievements* (Vendler 1957). Complex events may be analyzed as *causative*, linking two simple structures by means of a CAUSE relation, as in (2). The transitive verb *open* could then be given the decompositional representation in (3).[2]

(1) Simple event schema:
  a. $[\, x\ \mathsf{ACT}_{\langle \mathit{MANNER} \rangle}\ (y)]$
  b. $[\, x\ \langle \mathit{STATE} \rangle\,]$
  c. $[\ \mathsf{BECOME}\ [\, x\ \langle \mathit{STATE} \rangle\,]\,]$

(2) Complex event scheme
  $[\, [\, x\ \mathsf{ACT}_{(\langle \mathit{MANNER} \rangle)}\,]\ \ \mathsf{CAUSE}\, [\, \mathsf{BECOME}\, [\, y\ \langle \mathit{RES\text{-}STATE} \rangle\,]\,]\,]$

(3) *open*: $[\, [\, x\ \mathsf{ACT}\,]\, \mathsf{CAUSE}\, [\, \mathsf{BECOME}\, [\, y\ \langle \mathit{OPEN} \rangle\,]\,]\,]$

The complex causal schema in (3) has three important features:

(A) Only the second argument of the CAUSE relation (i.e., the result event) contains an idiosyncratic component.[3] The participant argument of the first event interacts solely with the primitive predicate ACT. This aligns well with the apparent non-specificity of an implied causing action for change-of-state verbs such as *open*: for instance, capturing the intuition that (3) expresses its subject's responsibility for the specified result state without constraining the method in which x acted on y to bring this state about.

(B) Per the inclusion of CAUSE, some analysis of causal relationships must underpin the connection between the argument of ACT and the predicate's result state. Neither (2) nor (3) commits to the nature of the causal relation, taking CAUSE as a (lexical) semantic primitive.

(C) The inclusion of both CAUSE and an associated result state in (2)/(3) gives rise to what we call the **culmination assumption:** that the denotation of an uninflected telic predicate exclusively contains eventualities which describe a complete developmental trajectory from

[2]Primitive (non-idiosyncratic) predicates are indicated in (1)-(3) without italicization; roots and ontological types are shown in angle brackets and italicized. Roots can occur as modifiers of primitive predicates (indicated in subscript, following Levin and Rappaport Hovav 2005) or as arguments. All optional elements are indicated parenthetically.

[3]Causative *change-of-state* verbs are typically taken to have decompositions of the form in (3); as (2) suggests, accomplishment predicates more generally may encode idiosyncratic manners in their process components (e.g., *walk home*). This issue is linked to the *Manner/Result Complementarity Hypothesis*; whether manner can be part of the first event in such decompositions depends on one's view of the MRCH, and lies beyond the scope of our discussion (Levin and Rappaport Hovav, 1991, 2013, 2014; Beavers and Koontz-Garboden, 2012).

some point of initiation up to and including the realization of the result state—i.e., a point of culmination, beyond which the event cannot continue.[4]

Complex predicates of the sort schematized in (2) are often categorized as **durative telic predicates** (Vendlerian *accomplishments*), in view of their association with a *process* (leading from initiation to termination) as well as a point of *culmination*. These culmination points represent categorical transitions between the developing eventuality and the onset of the (root-specified) result state. The outcomes or products of such culminations are thus precisely delineated by the decompositional representation of the underlying verb, and can range from the creation or destruction of an object (see 4a, 4b, respectively) to arrival at the terminus of a path (4c), to an instantaneous transition between states (of some pre-existing object, as in 4d).

(4) a. Emanuel baked a cake.
result: A (baked) cake. *object creation*

b. Maya ate a cookie.
result: A consumed cookie. *object destruction*

c. Des ran a marathon.
result: A complete (running) traversal of a 26.2 mile path *path termination*

d. Liina opened the door.
result: An open door. *state transition*

Prima facie, the schema in (2)—and the culmination assumption in particular—are in tension with the crosslinguistically robust observation that accomplishment predicates can be felicitously used to describe non-culminating eventualities. Contexts permitting non-culminating instantiation include, for instance, the complements of certain *aspectual verbs*, as shown for English in (5):

(5) a. Emanuel began to bake a cake.

b. Emanuel continued to bake a cake.

c. Emanuel stopped baking a cake.

Two things are immediately clear. First, none of (5a)-(5c) depend for their truth on the eventual existence of a complete (baked) cake. (5a), for instance, felicitously describes a situation in which Emanuel mixed together butter and sugar with the intention of making a cake, but went on to discover that the butter was rancid, putting a stop to the endeavour; (5b) might describe a situation in which he proceeded from some set of activities satisfying (5a) to a more advanced stage of the baking procedure (e.g, by adding flour and eggs to the butter-sugar mixture), but is compatible with his then being called permanently away for some emergency. Similarly, (5c) might narratively serve to introduce the emergency, as long as Emanuel's prior activities satisfy at least (5a).

Second, despite the apparent irrelevance of culmination, truth conditional assessment of (5a)-(5c) is linked to some relationship between Emanuel's reference time actions and the culmination specified by the embedded predicate. This relationship is what binds situations verifying (5a)-(5c) together as (partial or incomplete) instances of cake-baking, and what distinguishes these situations

[4] Anticipating the upcoming discussion, this point parallels Dowty's (1979, p.133) framing of the imperfective paradox: "the meaning of an accomplishment verb phrase invariably involves the coming about of a particular state of affairs. For example, *drawing a circle* involves the coming into existence of a circle [...], *kicking the door open* involves the door's coming to be open [...] The analysis of accomplishments in terms of BECOME-sentences was motivated [...] by the need to capture such entailments. Yet it is just this entailment that such a result-state comes about that fails when the accomplishment verb phrase appears in a progressive tense."

from activities like wiping counters, answering the phone, or chopping vegetables. Intuitively, situations validating (5a)-(5c) must plausibly be *parts* of a recognizable cake-baking process: that is, they must resemble the activities one might describe when asked to explain how to bake a cake.

The trouble is this: nothing in a decompositional representation like (2) supplies the sort of rich procedural information required to cash out the descriptive generalization above. While the acceptability of (5a)-(5c) is of course partially governed by the lexical semantics of the embedding verbs (*begin* selects for initiating activities, *continue* for intermediate stages, and so on), these embedding predicates must operate on the bipartite causal structure in (2). The 'procedural' part of a complex eventuality is specified in terms of structural (non-idiosyncratic) components: ACT (and its arguments) and CAUSE. As a result, any characterization of qualifying ACT eventualities can only be constrained by their causal relation to the result state, namely, by the standard assumption that causal (CAUSE) relations are realized only when the cause and effect both obtain. When an effect is not fully realized, this causal connection cannot be (fully) manifest, leaving the set of situations that might plausibly verify claims like (5a)–(5c) apparently unconstrained.[5]

This is evidently untenable. It is clear that certain activities (mixing butter and sugar together) are potential truth-makers for (5a)-(5c), while others (taking a phone call) are not. It is equally clear, moreover, that the information that determines whether or not a particular set of activities 'counts' as any part of a cake-baking is introduced by the predicate itself (and not by the embedding context). The preceding discussion thus illustrates an important point: truth conditional assessment of (non-culminating) uses of telic predicates relies on the availability of a rich body of *world knowledge* pertaining to the typical processes and procedures that bring about the appropriate culmination (and/or result state). While this information is evidently tightly linked to the specified culmination, it is—equally evidently—not itself made explicit by lexical means. The resulting puzzle of telic predicates is twofold. First, how is the world knowledge relevant for linguistic assessment mediated through the specified culmination? Second, by what formal means does this information enter into truth-conditional assessment? How these questions are answered depends on how the notion of a *process* leading to a certain result is to be captured and, consequently, on how *progress towards culmination* is to be recognized and formally assessed.

This paper aims to address the above challenges by introducing a new approach to the semantic structure of telic predicates, challenging the type of decompositional analysis presented in (2). We critically (re)evaluate implications (A)-(C), making the following claims.

(I) Against the (widely-accepted) culmination assumption (see C above), we argue that both culminating and non-culminating eventualities can validly instantiate telic predicates, as long as they exhibit an appropriate formal alignment with recognizable procedures for culmination.

(II) The representation of a telic predicate P does not hinge on a culmination *requirement*. It centers instead on a **culmination condition:** a specification of what can or must be true for an event of type P to culminate (cf. Kratzer 2004). Thus, unlike the decompositional scheme in 2 (see also A above), the predicate must—in addition to specifying a relevant result state—provide idiosyncratic information concerning what leads to this result. A telic predicate thus feeds truth conditional evaluation a structured body of information about the *causes* of culmination. This approach requires an account of causal relationships that goes

[5]The problem is specific to complex eventualities, due to the structure of their proposed decompositions. Per (1), predicates denoting simple events allow interaction between the idiosyncratic root and semantic primitives, so that contexts which reference the initiation or development of a simplex predicate have access to idiosyncratic information.

beyond a binary primitive such as CAUSE, which relates two eventualities by treating one as the cause and the other as the effect (see B above).

Concretely, we propose that the event type associated with a telic predicate P corresponds to a formal **causal model** (cf. Pearl 2000) in which P's culmination condition, C_P, occurs as a dependent variable. In the abstract, such a model comprises a set of **causal pathways** (informally: procedures; formally: sets of collectively sufficient causing conditions) for the realization of C_P, as well as pathways for its failure. The precise content of these pathways is provided by observation and/or experience, so that an event type model imposes set-theoretic structure on the world knowledge that constitutes our sense of how to bring C_P about: that is, on the sort of world knowledge that is crucial for truth-conditional assessment of (non-)culminating instantiations of P.

We wish to be clear at the outset about our analytic goals and the limitations of this paper. We aim to advance causal models as cognitive objects that bridge the gaps between lexically-specified information (i.e., information associated with particular expressions in the mental lexicon), world knowledge, and the set-theoretic objects on which a compositional semantic system relies. The causal approach is motivated for telic predicates by means of a well-known case study: Dowty's (1979) **imperfective paradox**. By revising the decompositional representation in (2)-(3), and letting go of the associated culmination assumption, our approach obviates the paradox, allowing us to assign an intuitive *partitive* semantics to the progressive operator. This semantics—presented in Section 4—is crucially a semantics for a progressive which composes with eventuality predicates analyzed as causal models: we associate progressive claims with formal conditions under which a causal model corresponding to the input event predicate can be said to be partially instantiated.

What we do not provide is a full account of the process by which a given eventuality predicate comes to be associated with a particular causal model in the mind of a speaker. Some aspects of this cognitive association must of course be experiential, but we acknowledge that others must be systematic, insofar as the aspectual class properties of certain eventuality predicates are known to emerge compositionally (Verkuyl, 1972, a.o.). It is our view that a theory of event structure based on causal models should ultimately amount to a new *aspect calculus* (cf. Dowty 1979), and should thus link the linguistic process of aspectual composition to the cognitive process of model construction, in such a way that familiar aspectual class properties can be "read" from the structure of the model. We see this paper (together with other recent work investigating the relationship between causal models and causal predicates; see Section 3) as offering a starting point for developing a causal theory of event structure, and hope to provide strong motivation for future work along these lines by showing that linking eventuality predicates to causal models can be instrumental in solving longstanding puzzles about the use and interpretation of these predicates.

The paper is structured as follows. Section 2 provides background on telic non-culmination in progressive contexts, focusing on data that pose difficulties for existing solutions to the imperfective paradox. Section 3 introduces *structural equation causal models* and shows how they can be used to explicate linguistic judgments. We offer a formal definition of telic predicates within this framework and demonstrate how the notion of process can be captured in terms of structured and multifaceted causal dependencies, rather than by means of a binary primitive such as CAUSE. Section 4 presents and illustrates a compositional treatment of the semantics of telic progressives. Section 5 returns to empirical and theoretical puzzles discussed in the prior literature, showing how the causal approach resolves some key challenges associated with the imperfective paradox; Section 6 explores further consequences of the analysis. Section 7 concludes and reflects on broader implications.

2 Culmination, expectations, and partial realization

2.1 Background: the imperfective paradox

The lexical puzzle of telic predicates has primarily been viewed through the lens of the **imperfective paradox**, which centers on the interpretation of accomplishment predicates under progressive modification (Dowty 1979). Unlike atelic predicates, which exhibit a symmetric pattern of entailment between their progressive and nonprogressive forms (as shown in 6), telic progressives do not entail their nonprogressive counterparts (Kenny, 1963).

(6) Atelic (activity) predicate:
  a. *Simple past (perfective):* Lila pushed a cart. $\leftrightarrow$
  b. *Progressive:* Lila was pushing a cart.

(7) Telic (accomplishment) predicate:
  a. *Simple past (perfective):* Lila drew a circle. $\rightarrow$
  b. *Progressive:* Lila was drawing a circle.

Intuitively, the lack of entailment from (7b) to (7a) is due to the distinct relationship that progressive and nonprogressive aspects have to the culmination condition of an uninflected accomplishment predicate. (7a) presents a completed picture, conveying that culmination occurred within reference time, and thus entailing the existence of a complete culmination product (i.e., a full circle, drawn by Lila). Progressive (7b) has no such requirement: like the aspectual verb constructions in (5a)-(5b), it requires only that some part of a circle-drawing coincides with reference time, so that Lila's actions *make progress* towards culmination but need not get all the way there.

Despite a clear contrast with (6), it is not immediately obvious that the asymmetric entailment pattern in (7) constitutes a 'paradox'. As formulated by Dowty (1977, 1979), the puzzle emerges from a clash between the empirical acceptability of telic progressives in contexts which preclude culmination, and standard assumptions about (a) the denotation of telic predicates, and (b) the contribution of grammatical aspects. The prevailing **culmination assumption** is that the complex eventualities denoted by telic predicates necessarily incorporate the point of culmination, a perspective that gains support from the truth conditions of telic perfectives in many languages (exemplified for English by the simple past in 7a), which license **culmination entailments**.

Culmination entailments are readily explained by combining the culmination assumption with an extensional treatment of the perfective aspect (see 8), on which the 'completed-event' perspective is captured by including the runtime ($\tau(e)$) of an instantiated eventuality (e) in the reference time (t) supplied by tense.[6] Since, per the culmination assumption, e qualifies as a P-event just in case it includes a realization of P's culmination condition (C_P), perfective modification forces a reference time realization of C_P. The culmination entailment follows immediately:

(8) $[\![\textsc{pfv}]\!] := \lambda w \lambda t \lambda P_{vt}.\exists e[\tau(e) \subseteq t \wedge P(e)(w)]$ (Klein, 1994; Kratzer, 1998)

(9) $[\![(7a)]\!]^{w^*,t^*} = [\![\textsc{pst}(\textsc{pfv}(\sqrt{}\text{Lila draw a circle}))]\!]^{w^*,t^*}$
$= \exists e[\tau(e) \subseteq t^*\{\prec \mathsf{now}\} \wedge \mathsf{draw\text{-}circle}(\mathsf{L})(e)(w^*)]$

[6]This formulation assumes a three-times system (Reichenbach, 1947; Klein, 1994), on which aspects relate event time to reference time, and pronominal tenses (Partee 1973; Kratzer 1998) establish a relationship between reference time and speech time. Events (type v) and times (type i) are related via Krifka's (1989) *temporal trace* function, $\tau : D_v \rightarrow D_i$, with $\tau(e)$ returning the runtime of an input eventuality e.

In contrast to the completed perspective of the perfective, progressives portray a given situation as *ongoing* (in progress) at reference time. Extrapolating from above, then, a natural semantics for the progressive would reverse the key inclusion relation in (8) to yield (10).

(10) $[\![\text{PROG}]\!] := \lambda w \lambda t \lambda P_{vt}.\exists e[\tau(e) \supset t \wedge P(e)(w)]$

(11) $[\![(7b)]\!]^{w^*,t^*} = [\![\text{PST}(\text{PROG}(\text{Lila draw a circle}))]\!]^{w^*,t^*}$
$= \exists e[\tau(e) \supset t^*\{\prec \mathsf{now}\} \wedge \mathsf{draw\text{-}circle(L)}(e)(w^*)]$
entails: $\exists t'[t' \supset \tau(e) \wedge \mathsf{draw\text{-}circle(L)}(e)(w^*)]$

Combining the culmination assumption with (10) thus straightforwardly predicts that the progressive of a telic predicate P can only be true if reference time events continue to develop to culmination. But this goes against the empirical data: as demonstrated by (12), telic progressives can be both acceptable and true in contexts which rule out any possibility of culmination.

(12) Mahler was composing his tenth symphony (when he died). (True)

Since Dowty's 'paradox' follows from combining the culmination assumption with an extensional treatment of (the progressive) aspect, the two obvious routes towards its resolution involve modifying one of these assumptions. Both approaches—revising the extensional progressive in (10) and relaxing the culmination assumption—have been explored in the literature, and we summarize the key ideas of each approach below.

(I) **Type I: Intensional progressive approaches**
The culmination assumption can be preserved by intensionalizing the progressive operator (Dowty, 1979; Landman, 1992; Asher, 1992; Bonomi, 1997; Zucchi, 1999, a.o.), permitting it to instantiate culminated P-eventualities in alternatives to the evaluation world. On this approach, PROG(P) is true of situation s at time t in world w just in case s continues to develop to the realization of C_P in some set of worlds sufficiently similar to w at time t.

The key analytical challenge is to appropriately constrain the alternatives introduced by PROG: this set must plausibly exclude the evaluation world to allow for data like (12). Simultaneously, however, the alternatives must track actual events fairly closely, in order to capture the intuition that reference time facts correspond to some initial portion (or *stage*; Landman) of a P-event—i.e., to establish that a P-event is actually *in progress* in w at reference time.

(II) **Type P: Extensional or partitive approaches**
The other option is to discard the culmination assumption, allowing (non-)culminated eventualities to validly instantiate a telic predicate (Parsons 1989, 1990; see also Lascarides 1991; Szabó 2004, 2008). In principle, this move obviates the imperfective paradox, allowing us to account for data like (12) while maintaining an extensional treatment of the progressive.

The analytical challenge here is to establish what constitutes a valid *partial realization* of a telic predicate. In the context of (12), for instance, a satisfactory account must explain what attributes of an unfinished symphony-writing activity align it with a completed instance of the same predicate, while distinguishing it from seemingly comparable occurrences of a related but atelic activity (e.g.,*composing music*).[7]

[7] From the Type P perspective, the imperfective paradox is more appropriately termed the *partitive puzzle* (Bach 1986); its solution extends beyond a treatment of telic progressives and promises to explain the interpretation of telic predicates in other non-culminating contexts (such as those exemplified in 5).

Our view is that the two analytical challenges outlined above are intertwined. Whether one adopts an extensional or intensional analysis of the progressive operator, truth conditional assessment of 'paradoxical' data hinges on the concept of *partial realization*. This notion must be a partitive one, insofar as reference time facts can only constitute a partial (non-culminated, process-only) P-type event if there is some abstract or concrete whole of which these facts form a well-defined part. At the same time, partial realization (at least in the progressive context) must have an intensional quality: an *in-progress* P-event must naturally reflect a progression of facts that has advanced some way towards a (not yet realized but in some sense realizable) culmination condition.[8]

Looking ahead, we use causal models to combine partitive and intensional perspectives in a general theory of telic predicates. This will enable us to provide a clear definition of the abstract, culmination-linked whole of which an in-progress P-eventuality must be a component. Unlike standard Type I analyses—which introduce intensionality exclusively via the progressive operator—the causal-modeling approach builds the crucial intensional component of the imperfective paradox directly into the denotation of a telic predicate. As a result, we avoid making the 'partiality' of reference time activities contingent on the locally-projected likelihood of culmination. This will allow us to explain some key data that has persisted as a challenge to the received Type I approach.

Before introducing our formal proposal, we review some key developments in Type I approaches, focusing our discussion on the challenges posed by three types of data: progressives associated with **impossible events**, **unlikely events**, and scenarios in which culmination is locally but not categorically **out of reach**. We use these data to lay the groundwork for an approach that integrates a guiding principle of the Type P approach: namely, that reference time facts matter more than modal projections in determining whether or not a given situation constitutes a valid partial realization of a telic predicate. We conclude this section with a new descriptive generalization for the truth conditions of telic progressives, which is the basis for our formal account.

2.2 Inertia and expectation

On the Type I view, *partial realizations* of telic predicate P 'naturally' lead to the realization of P's culmination condition C_P. $\textsc{prog}(P)$ is true of situation s just in case s continues to develop to the realization of C_P in a set of normal alternatives to the evaluation world. This is intuitively appealing insofar as it captures the idea that an in-progress (process-only) P-eventuality 'aims' at culmination. However, it has proven difficult to characterize the relevant alternatives in a way that properly distinguishes acceptable and unacceptable uses of telic progressives.

Dowty's influential (1979) analysis aims to cash out an appropriate modal relationship in terms of *inertia*: per (13), reference time occurrences verify $\textsc{prog}(P)$ just in case they eventually realize C_P if allowed to continue uninterrupted. The set of *inertia worlds* $\mathsf{INR}(w, t)$ projected from w at time t are defined as those worlds which share a history (Thomason 1984) with w through t, and which are further constrained to develop after t in the way that is "most compatible with the past course of events" (Dowty, p.148): the result is that a (future) realization of C_P is treated as the normal outcome of world history as considered from the reference time perspective.[9]

[8] To the extent that an in-progress or non-culminated instantiation of a telic predicate P describes a *process* component of P, the tension between Type I and P approaches (and between different theories of each type) hinges on how the notion of a process for culmination is understood. Although we do not frame our discussion of progressives in terms of the process concept, our proposal takes a clear position on this question; see Section 3.5.

[9] Compatibility is based on Lewis's (1973b; 1979) *similarity*, which is related to *stereotypicality* (Kratzer, 1981). The type of information needed to assess stereotypical development—that is, that guides expectations about the

(13) **The inertia world analysis** Dowty (1979, p.149)
For a telic predicate P, $\textsc{prog}(P)$ is true at $\langle w, t\rangle$ iff, for some interval t' such that t is a nonfinal subinterval of t', and for all worlds $w' \in \mathsf{INR}(w,t)$, P is true at $\langle w', t'\rangle$.

Dowty's reliance on a global perspective (i.e., one which takes an entire world into account) turns out to have undesirable consequences (Vlach, 1981): linking inertia worlds to the complete history of the evaluation world incorrectly predicts falsity for telic progressives whenever an event that actually blocks culmination is foreseeable at reference time. This makes the wrong prediction for (14): since $\mathsf{INR}(w,t)$ picks out worlds in which all reference time a activities continue as established, it includes worlds which contain both the inertial continuation of Henrietta's trajectory, as well as that of the offending truck. As a result, the fatal collision must occur across $\mathsf{INR}(w,t)$. This should immediately falsify (14a), but the claim is judged to be both acceptable and true in context.

(14) **The collision scenario.** Henrietta walked into a crosswalk, intending to cross to the other side of the street. As she took her first step into the street (but unknown to her), a truck was racing towards the same crosswalk with a speed and trajectory that placed it on a collision course with Henrietta (given her speed and trajectory).

a. *Target:* Henrietta was crossing the street (when the truck hit her).

The problem can be avoided by refining Dowty's proposal so that the information used to project inertial futures is a proper part of the global picture. One way of doing this is to allow the inertial development of only those processes which are understood to belong (in some sense) to the target predicate (cf. Landman 1992; Bonomi 1997; see Section 5.2).[10] Alternatively, the same basic intuition can be captured by restricting attention to a particular *perspective* on the reference time situation, roughly constituted by a subset of the (global) facts at reference time (Asher, 1992).

Asher's account is technically complex, but its central idea can be expressed fairly simply. Intuitions about the role of normal or inertial developments are expressed in terms of *default* inferences, concerning what typically follows from a particular body of information. On this view, the progressive of a telic predicate P can hold at time t just in case there is some *admissable perspective* π on the full set of reference time facts (i.e., some suitably-constrained subset of these facts) such that all of the maximally normal worlds in which π holds are worlds in which the nonprogressive form of P holds at or after reference time (s normally leads to C_P).

In addition to explaining why a telic progressive can hold in 'predictable failure' contexts such as (14)—the perspective adopted here is ostensibly that of Henrietta, who is not aware of the oncoming

normal developments of a situation—is very plausibly causal in nature (see, e.g., Kaufmann 2013; Nadathur 2023a).

[10] Landman (1992) calls this the *normality proposal*; the idea is to relativize inertia worlds to an event e as well as a world-time pair, so that world w differs from an inertial alternatives w' at most in that e follows its 'natural course' in w'. This differs from Vlach's (1981) solution insofar as Vlach links the truth of a progressive to the *causal consequences* of reference time events, rather than their normal developments. Landman finds Vlach to be insufficiently specific about the parameters under which events in progress are assumed to continue, pointing out that if (14a) is used at the moment that Henrietta steps into the street, the evaluation world is one in which her crossing continues past reference time, but is not a world in which this process leads to the truth of the corresponding non-progressive. Thus, for Landman, the proposal must be refined to require that the ongoing process continues *past the point at which it is interrupted in* w^*, and this forms part of the basis for his own proposal (see Section 5.2).

The causal component of Vlach's proposal is compatible with our view. However, it remains subject to an objection raised by Bonomi (1997): the mere existence of a causal relationship between reference time activities and the culmination condition of predicate P cannot license $\textsc{prog}(P)$. Henrietta's waking up in the morning may eventually lead to her crossing the street, but neither (14a) nor its present tense correlate can be valid before she first steps into the street. Thus, while causal relationships between reference time actions and C_P are truth-conditionally relevant for $\textsc{prog}(P)$, a causal analysis is incomplete without an accompanying mereological view of the target event type.

truck—Asher's restricted normality-based approach explains the observation that progressives of telic predicates with (locally) contradictory culmination conditions can be true of the same general state of affairs. Both (14a) and (15) are truthful descriptions of the context in (14), even though worlds in which the culmination associated with (15) occurs will be worlds in which Henrietta does not reach the other side of the street (Abusch, 1985).

(15) Henrietta was walking to her death.

Asher's explanation is straightforward: "it suffices for the truth of the progressive that there be just one perspective π on [a state of affairs] s such that the normal course of events based on having a state with the characteristics given in π leads to a completion of the appropriate kind" (1992, p.480). Thus (14a) can be true because s contains a set of facts (pertaining only to Henrietta and her activities) that licenses a default inference to an eventual successful crossing; however, since s also contains a larger set of facts including the truck, there is an alternative perspective on the situation licensing a default inference to the fatal collision.

Broadly then, Type I approaches tie the truth conditions of telic progressives to the validity of some (local) expectation of culmination, while allowing this expectation to be based on what may be an extremely narrow construal of the evaluation world facts (see also Lascarides 1991 on the "eventual outcome strategy"). This seems to capture the observed unacceptability of progressives of **impossible tasks**, such as those in (16):

(16) a. #/?Mary is swimming across the Atlantic Ocean.
b. #/?The children are digging a hole to China.

It is certainly possible for Mary to get into the ocean at Galway and begin to swim westward, but there is no realistic perspective which admits both basic human capabilities and the physical properties of the Atlantic while simultaneously predicting successful crossing as the typical outcome of her reference time actions. Similarly, for (16b), while it might be true that some children are digging a hole in the sand at place on the Earth's surface which is antipodal to China, no perspective which admits basic facts about the earth's composition, physics, and the limits of human strength and heat tolerance will license a default inference to the success of an attempt to dig all the way through. On a normality-based approach, then, the claims in (16) are false: their target outcomes fail in all normal projections based on (rational subsets of) reference time facts.

However, the same prediction of falsity incorrectly applies to progressives of tasks which are possible, but extremely unlikely to be completed.[11] (17) illustrates: despite the lack of a minefield-aware perspective licensing a default inference from an attempted to a successful crossing, (17) can be true when Trapper John has taken some steps into the danger zone.[12]

(17) Trapper John was crossing a live minefield.

A similar example from Bonomi (1997) shows that the progressive of telic P can be true even where non-culmination (the failure of C_P) is the most likely outcome of any good-faith attempt at P. The

[11] By *impossible events* we do not mean only logically impossible events (as in *Mary is drawing a square heptagon* or *Mary is proving that* $2+2=5$), nor is the distinction between impossible events and unlikely events intended as a scalar one, as though impossible events were simply extremely unlikely. Rather, the relevant cases of impossible events involve culminations that are incompatible with the physical constraints ordinarily presupposed in interpretation. For example, swimming across the Atlantic and digging a hole to China (see 16) do not represent low-probability events, but describe outcomes that are, under normal physical assumptions, impossible for a human to bring about.

[12] Note that the acceptability of (17) does not rely on an "ignorant" perspective: it remains valid in a context where all participants, Trapper John included, are apprised of the danger (M*A*S*H, Season 2, Ep.6).

context for (18) and the spokesman's presumed expertise block a default inference to the success of any single boat in circumnavigating the globe; (18a) is nevertheless both acceptable and true.[13]

(18) **The sailing scenario.** Suppose that an international sailing association organizes a competition whose goal is the circumnavigation of the globe. After thorough selection, 100 boats are admitted, and all sail on the scheduled date. Some days later, a spokesman says:

a. 100 boats are circumnavigating the globe. Most of them will fail.

From one point of view, the difference between unacceptable progressives of impossible events and potentially-valid progressives of merely **unlikely events** has to do with the set of possible outcomes of an attempt to do P. Both Dowty and Asher treat the progressive as a type of *universal* modal, requiring successful culmination across the full set of relevant alternatives. Impossible task progressives fail this requirement because *none* of the normal/inertial alternatives (whether generated globally or from a restricted perspective) realize C_P, but the same need not be true of unlikely event progressives. We might then thread the needle between these two types of data by weakening the modal force of an intensional progressive operator: the progressive of some predicate P would then be true of a reference time state s just in case some admissable perspective on s has at least one normal continuation that realizes C_P.

As it turns out, even an existential intensional progressive makes the wrong predictions for telic predicates in contexts where culmination is *locally* **out of reach**.[14] The target task in (19), for example, is not categorically impossible, but successful culmination is contextually precluded by (potentially mutable) properties of the protagonist Benny:[15]

(19) **The undertrained runner** (modified from Varasdi 2014, pp.192–193). Benny is an amateur but ambitious runner who decides to run an ultramarathon. He will not in fact be able to run the whole distance, because very few can, and he does not have the necessary physical qualities, having trained insufficiently to build up his endurance to the required degree. Nevertheless, he registers for the race, appears on the given date, gets his race number, and sets off with the other runners. The first couple of kilometers go well, but then he begins to lose strength, and at about half the distance he collapses because of exhaustion.

a. *Target:* Benny was running an ultramarathon (when he collapsed).

From an 'objective' perspective—that of a speaker who has learned Benny's name, but does not know him—the problem with (19a) is the same as (18). We may be able to define a perspective

[13] As Bonomi points out, Asher's truth conditions for (18a) require that (i) is a valid default; the spokesman's first and second sentences should be contradictory, but they appear perfectly compatible.

(i) For any boat x which is circumnavigating the globe, x typically circumnavigates the globe, eventually.

[14] Dowty (1979) (citing R. Thomason) sets aside existential analyses of the progressive based on (ii): since a fair coin will come up heads in half of the inertial futures, (ii-a) and (ii-b) should both be true, against intuition.

(ii) **The fair coin scenario.** Assume a fair coin has been tossed, and at reference time is still rising.

a. ??The coin is coming up heads.
b. ??The coin is coming up tails.

[15] Szabó (2004) provides example (iii), illustrating much the same point as (19):

(iii) As the architect was building the cathedral he knew that, although he would be building it for another year or so, he couldn't possibly complete it.

on (19) which excludes details about Benny's training and endurance, but the default inference required for the truth of (19a) (namely, that completing an ultramarathon is the *typical* result of running in one) should still be precluded by domain-relevant information: most attempts to run an ultramarathon do not conclude successfully. The problem goes beyond this, however: (19a) seems to be acceptable and true even when we assume it to come from a 'subjective' speaker (say, a close friend), who is aware of Benny's physical limitations and thus the inevitability of his failure.

The judgments for (19a) in an out of reach context simply cannot be explained on an inertia- or normality-based approach to the progressive, regardless of modal force: *all* of the normal futures projected from (a subjective view of) the reference time situation must be ones in which Benny does not complete the race. If the intensional approach is correct, out of reach progressives should pattern with impossible event progressives. That this is evidently not the case shows clearly that there is some distinction between impossible and out-of-reach tasks to which the progressive is sensitive. The received Type I approaches fail to account for this distinction.

The intended contrast between *unlikely events* and *out-of-reach events* is not simply between two kinds of difficult events, but between events that are unlikely in a general sense and events whose probability is effectively zero in the specific context at issue. In unlikely cases such as (18), the context states that "most of [the 100 boats circumnavigating the globe] will fail," so a situation in which all 100 succeed is unexpected, but not blocked by any principled contextual constraint. In out-of-reach cases like (19), by contrast, the relevant preparatory conditions are absent: someone who has not completed the necessary training cannot run an ultramarathon. Such events are effectively impossible in the given circumstances, though not impossible for humans as such. This differs from cases like (16), which describe tasks that are generally impossible for a human agent.

The preceding discussion suggests that an account based on completion requirements in all or even some 'normal' worlds is not the right approach for telic progressives. Indeed, Szabó (2004, 2008) takes data like (19) (see also fn.15) as evidence that the truth of telic progressives is not grounded in the modal accessibility of culmination: "events that cannot possibly culminate can nevertheless be truly said to be in progress, as long as we are allowed to abstract from those features that made their completions impossible" (2004, pp.38–39, see also Varasdi 2014). From this viewpoint, the right approach must be a partitive one, on which the truth of telic progressives hinges on what is actually going on at reference time, rather than the 'normal' or expected consequences of these events. This is consistent with what Landman (1992, p.13) calls the 'part-of proposal'.

(20) **The part-of proposal.** *Mary is crossing the street* is true iff some actual event realizes sufficiently much of the type of events of Mary's crossing the street.

The load-bearing element in (20) is the phrase "sufficiently much". A partitive approach to the imperfective paradox must make reference to some set of criteria according to which an in-progress but non-culminated eventuality 'counts' as (part of) a P-event. This set of features must, intuitively, have a connection to the culmination condition of the underlying predicate P. However, if the partitive approach is to fare better than the intensional progressive on 'out-of-reach' data like (19), the connection cannot be formalized purely in terms of a local (modal) expectation of culmination. Put another way, any relationship between reference time activities and P's culmination condition must be realized at the level of the abstract *event type* denoted by P, rather than in terms of a reference time token. The challenge thus lies in establishing precisely what it means for some set of (reference time) facts to realize *sufficiently much* of a P-type event.

2.3 Parts and measurement

Parsons (1989, 1990) develops a well-known extensional (Type P) account of the imperfective paradox, based on the idea that in evaluating telic progressives, "present activities [...] are the whole story," rather than "the idea of what would be the case [...] if present activities were to go on uninterrupted" (1989, pp.221-222). This approach necessitates letting go of the culmination assumption, so that non-culminated eventualities can provide valid instantiations for telic predicates.

Parsons shifts responsibility for culmination entailments (where they arise, as in 7a) from the denotation of an uninflected telic predicate to the semantic contribution made by a modifying aspectual operator. In his framework, progressives and nonprogressives do not differ in their intensionality, but in which features (how much) of a P-type eventuality they require reference time events to realize.[16] Example (21) schematizes Parsons's proposals: both perfective (21a) and progressive (21b) require that a reference time event e satisfies some core characterization provided by the uninflected predicate, but differ in whether or not e must *culminate* at reference time, or can merely *hold*. Crucially, while the relation $\mathsf{Cul}(e, t)$ in (21a) is true of e at time t just in case e culminates at t, $\mathsf{Hold}(e, t)$—imposed by progressive (21b) instead of Cul—neither implies nor predicts culmination, requiring only that "e is an event which is in development at t" (1989, p.220).

(21) a. Mahler wrote a symphony. $\sim$ PST(PFV($\sqrt{}$*Mahler write a symphony*))
$\exists e[\mathsf{writing}(e) \wedge \mathsf{Agent}(e, \mathsf{m}) \wedge \mathsf{Theme}(e, \text{a symphony}) \wedge \mathsf{Cul}(e, t\{\prec_i \mathsf{now}\})]$

b. Mahler was writing a symphony $\sim$ PST(PROG($\sqrt{}$*Mahler write a symphony*)).
$\exists e[\mathsf{writing}(e) \wedge \mathsf{Agent}(e, \mathsf{m}) \wedge \mathsf{Theme}(e, \text{a symphony}) \wedge \mathsf{Hold}(e, t\{\prec_i \mathsf{now}\})]$

To the extent that this approach obviates the imperfective paradox, it has a clear conceptual appeal. Analytically, however, Parsons's explication of Hold leaves something to be desired: namely, what it means for an event in the denotation of a telic predicate to be "in development". A partial answer emerges in the logical consequences of combining a Hold requirement with the first three conjuncts of (21b), which come from the (uninflected) sentence radical ($\sqrt{}$*Mahler write a symphony*). Specifying a semantic Theme in a representation like (21b) forces the reference time event e to involve some object belonging to the denotation of the nominal *symphony* (see also Landman, 1992). In conjunction with Hold, however (as opposed to Cul), this thematic object can at best be a *partial realization* of a symphony. Parsons sees this as a positive result: he accepts that the truth conditions in (21b) commit him to the existence of a symphonic object, but suggests that where event e holds (but does not culminate), the object is simply incomplete.

For Parsons, then, partitivity (partial realization) for telic predicates reduces to partitivity in the nominal domain (see Link 1983; Bach 1986 on parallels between the nominal mass/count distinction and the atelic/telic contrast that underlie this idea). On this view, incomplete objects can be considered instantiations of their intended wholes: Parsons suggests that non-culminated

[16] In the standard approach (see 8, 10), both perfective and imperfective aspects compose with predicates of events. Parsons adheres to this view. An alternative, metaphysical approach from the philosophical literature instead associates each aspect with a distinct type of ontological entity. Specifically, imperfective sentences assert the occurrence of a 'process,' while perfective claims assert that an 'event' has occurred (Mourelatos 1978; Stout 1997; Steward 2013; Baratella 2023.) The metaphysical perspective avoids the imperfective paradox by treating culminations events and thus not entailed by processes. According to Steward and Baratella, processes picked out the progressive have an intensional component involving possible worlds. However—unlike the account we develop here—these proposals do not address how to verify the instantiation of either a partial or complete process. These accounts also fail to explain why telic perfectives entail their progressive counterparts, as illustrated in (7). See Section 5 for further discussion, and Section 3.5 for an elaboration on the notion of process in different approaches.

events can similarly instantiate telic predicates, as their (in)completeness can be measured by the (in)completeness of the objects involved. However, this link between event and nominal partitivity relies on the existence of a tangible correlate of development and fails with accomplishment predicates that lack appropriate thematic arguments (in particular, *incremental themes*; Dowty 1991). Landman (1992) contends, for instance, that certain 'sudden' creation processes, as in the progressive *God was creating a unicorn*, involve necessary developmental stages that partially instantiate the telic predicate in the absence of even an incomplete physical object.

If judgments about divine creation seem obscure, Landman's argument can be reconstructed with a more mundane eventuality, such as baking a cake. This type of creation process is not sudden, but is crucially non-homogeneous: while a completed cake is a fully baked product, and there are various tangible correlates of development, no object that can be described as even an incomplete cake comes into being until the events in question are well underway. For instance, an event in which some agent mixes together appropriate ingredients intuitively counts as a cake-baking in progress, even though the contents of the mixing bowl cannot at this point be called a cake. This issue is even more apparent at earlier stages of cake-baking: *Emanuel is baking a cake* is true when he has merely collected the right ingredients and preheated the oven, but there is at this point no single object which could constitute the thematic cake required by Parsons.

We believe that there is something important in common between telic predicates of sudden or 'non-homogeneous' creation and those with incremental themes. In each case, the truth of the progressive hinges on some measure of progress through a sequence of steps which lead in principle to the right culmination. For predicates with incremental themes, it just so happens that steps along the pathway towards culmination are monotonically tracked by changes in the physical extent of an object, but this effect is secondary: it is not the presence of the physical correlate that validates a progressive claim. The cake-baking example underscores this point: even absent a(n incomplete) thematic object, Emanuel's actions make progress toward cake-baking: he has taken steps necessary for cake-baking that correspond, nearly literally, to the steps one might find in a cake-baking recipe. This correspondence allows us to evaluate both whether he is baking a cake at all as well as how far along he is in this process (*how much* of a cake-baking event he has realized).

The perspective outlined here aligns with Parsons in certain respects. First, it focuses on actual events (rather than projected future developments) as the entities that can verify or falsify a telic progressive. We also concur that event partitivity is conceptually parallel to nominal partitivity, and that a treatment of the former must incorporate a framework for measured comparison. Crucially, a partial realization of some event type (like a partial object) is partial because of the relationship it bears to a conceptual whole: it must contain at least some of the components that can be identified in the whole.[17] However, we do not believe that event partitivity can be assessed using the same measurement framework as object partitivity: the former does not reduce to the latter. Instead, the conceptual whole in the event domain must correspond to a well-defined 'recipe' or pathway for achieving culmination. The intuition is this: to decide whether Emanuel is baking a cake, one must first know *how a cake is baked*. Only then can one assess whether Emanuel's actions correspond, step by step, to the components of the conceptual whole that constitutes cake-baking.

Insofar as it requires a framework for measuring progress through a culmination 'recipe', partitivity in the domain of complex (telic) eventualities is dependent on the specification of a cul-

[17]Put another way, just as identifying an object as *part of* a house involves comparing its components to those of a complete house, identifying an actual event as *part of* a house-building involves comparing its components (subevents or stages) to a complete or *culminated* version of the same kind of thing.

mination condition. Our task, then, is to find some way of formalizing the relationship between process steps (line items in a recipe, to continue the metaphor) and culmination in such a way that measurement and comparison can be achieved on the basis of "present activities" alone. The link between a culmination procedure and the successful culmination of a telic predicate must therefore be established within the denotation of the uninflected predicate, and cannot rely on local (reference time) modal projections about the future. We will argue that causal information lies at the heart of the process-culmination relationship.

2.4 Taking stock: interim summary and desiderata

The preceding sections established desiderata for a satisfactory account of telic non-culmination in the progressive context. We review these points here, using them to motivate a particular descriptive generalization as the target of our analysis, and to clarify the analytical challenges which a causal-modelling approach is intended to address.

Section 2.2 introduced three key classes of data: (telic) progressives of *impossible events* (IEs), progressives of *unlikely* or *unexpected events* (UEs), and progressives in *out of reach contexts* (ORCs). Since intensional Type I approaches require reference time facts to predict successful culmination across an appropriately constrained and locally-projected set of modal alternatives, these accounts predict falsity across the board. This is compatible with empirical judgments for IE progressives (see 16a-16b), but it does not reflect the empirical data in the case of UE and ORC progressives, which can be both acceptable and true, as shown in Table 1.

	Examples	Empirically	Type I prediction
IEs	(16a) Mary is swimming across the Atlantic Ocean (16b) The children are digging a hole to China	X	false
UEs	(17) Trapper John was crossing a minefield (18) 100 boats are circumnavigating the globe	✓	false
ORCs	(19) Benny was running an ultramarathon	✓	false

Table 1: Intensional progressive (Type I) predictions for IE, UE, and ORC progressives

Based on the judgments reported above for UE and ORC progressives (see also data from from Bonomi 1997; Szabó 2004, 2008; Varasdi 2014), we argue that the central tenet of Type I approaches—i.e., that the truth of telic progressives relies on locally-projected alternatives—cannot provide an empirically adequate account. Appropriate truth and acceptability conditions must instead be based on something that the input predicates for UE progressives and those for ORC progressives have in common with one another, but fail to share with predicates of IEs. Moreover, since UE and ORC progressives do not share the same (local) intensional relationship to culmination, the pattern of judgments reported in Table 1 suggests that, where a telic progressive is acceptable, its truth conditions must depend on information pertaining to reference time facts, and not to their expected developments.

This suggests that a partitive Type P approach might be preferable; we therefore discussed Parsons's (1989; 1990) account in Section 2.3. While we agree with the spirit of his analysis—in particular, that a telic progressive requires the evaluation world instantiation of an eventuality constituting a well-defined *part* of a complete (culminated) eventuality denoted by telic predicate P—we argued that the requisite conception of *parthood* cannot be based on a tangible or physically measurable correlate of the progress of a P-eventuality. We suggested, instead, that the appropriate basis to use in comparing reference time facts to a complete P-eventuality is that of a procedure or 'recipe' for culmination: more precisely, that an evaluation world event qualifies as a potential *part* of a P-eventuality just in case it realizes some portion (set of steps) that belong to a recognized *procedure* for realizing P's culmination condition. With respect to the Table 1 data, this means that the sense in which Trapper John is crossing a minefield, the competitor boats are circumnavigating the globe, or Benny is running an ultramarathon at their respective reference times is that in each case, the (relevant) set of reference time facts corresponds to some part of a culmination procedure which one might adopt in setting out to complete an event of the right type.

The notion of a **culmination procedure**—to be given formal content below—not only accounts for the acceptability of UE and ORC progressives in the contexts summarized in Table 1, but also distinguishes them from the unacceptable IE progressives in (16a)–(16b). UE and ORC tasks allow progress toward culmination because there is a recognizable procedure by which culmination could be achieved. IE tasks, by contrast, are impossible because no such procedure is available, and these tasks therefore provide no basis for partitive assessment of a reference time event.

We propose the following descriptive generalization as our analytical target:

(22) **Descriptive truth and felicity conditions for telic progressives.** Let P be an accomplishment predicate with culmination condition C_P. Then:

a. $\textsc{prog}(P)$ is felicitous (hence eligible for truth-conditional evaluation) just in case P is associated with at least one culmination procedure; informally, a 'recipe' for C_P.[18]

b. Where felicitous, $\textsc{prog}(P)$ is true just in case the following conditions hold. First, the reference time situation s can be (strictly) partially matched to a culmination procedure associated with P—that is, s maps onto a snapshot or cross-section of a partial P-eventuality, defined in terms of partial realization of such a procedure. Second, some fact in this partially-realized culmination procedure has changed in value compared to the situation prior to reference time. Finally, the facts in s do not already guarantee that C_P will be false.[19]

Per (22), the truth and felicity of $\textsc{prog}(P)$ do not depend on locally projected continuations of the reference time situation. They do, however, require a formal relation between what is true at reference time and the structure of an abstract P-eventuality, itself defined in terms of culmination; the requirement that s realizes some change as compared to the pre-reference time situation simply establishes that something is in fact *happening* (has happened) at reference time.

The adequacy of this analysis depends on how the central notion of a *culmination procedure* is formally developed: Section 3 provides a framework for the proposal in (22). Our central idea is that a culmination procedure is defined in causal terms articulated at the *type* rather than the *token* level. We conclude the present section by sketching what such an approach amounts to, and by identifying the main challenges facing an adequate treatment of culmination procedures and their relation to the lexical representation of telic predicates.

[18] Lascarides (1991) suggests that the content of a *process eventuality* is supplied by world knowledge; we ultimately propose that world knowledge about such processes receives set-theoretic articulation *via* a formal causal model.

[19] Put another way, in order for an instantiation of P to be ongoing at time t, it must be possible at t to make additional progress towards the culmination condition C_P. Varasdi (2014) also incorporates such a requirement, by means of establishing that certain so-called "sustaining" conditions for a predicate must hold at the time of evaluation.

For current purposes, the distinction between *type*-level as opposed to *token*-level causation is closely tied to the nature of the relationship between partial and complete (culminated) instantiations of a complex predicate—or, more specifically, between a partial instantiation and the relevant culmination condition. As Szabó (2004) argues, the truth of a telic perfective requires not only the (prior) truth of its progressive correlate and the realization of the culmination condition of the underlying predicate, but also the actual instantiation of a causal relation between the progressive eventuality and this culmination.[20] However, since telic progressives can be true in contexts where culmination does not occur, it cannot be the case that an actual (in-progress) event can only verify the progressive in case it constitutes an *actual cause* for the relevant culmination condition.[21]

This is the central tension. On the one hand, the information contributed by a telic predicate must make it possible to verify some set of actual causal relations—either between a developing process stage of a P-eventuality and its realized culmination, or between developmental stages of an incomplete P-eventuality (see Szabó 2004 and fn.20)—*modulo* how much of a P-eventuality is in fact instantiated. On the other hand, making the truth of a telic progressive dependent on an actual causal relation to culmination precludes an adequate treatment of many valid uses of the progressive, including (for instance) ORC cases, since in such cases no culmination occurs. Hence, whatever qualifies a course of events as a valid partial instantiation of a telic predicate—or, in the terms of (22), whatever licenses a match between a reference time situation and a stage in a culmination procedure for P—must be identifiable independently of actual causation.

We therefore suggest that accomplishment predicates are not defined in terms of instances of actual causation, and that their inflected uses do not correspond directly to causative claims. Nevertheless, the information contained in a telic predicate is of the sort that *can* license actual causal claims, explaining (for instance) why culminated uses of such predicates entail actual causal links between some prior event and the realized culmination. We argue that singular causal claims (i.e., claims that a causal relation was realized at a particular moment) are licensed by **type-level causal information**, that is, by a body of knowledge comprising generalizations about causal relations in the world. Type-level information contrasts with **token-level information** (descriptions of actual causation) in that it specifies what *would be needed* for a particular result to

[20]On the basis of acceptable ORC progressives such as (19a) and (iii) (see fn.15), Szabó (2004, p.40) concludes that Type I approaches to the imperfective paradox are untenable. Rather than analyzing telic progressives in terms of their perfective correlates, Szabó proposes a 'reverse' analysis, on which telic perfectives entail strictly more than progressives (see 7). His proposal, sketched in (iv) below, makes clear not only that an actual causal relation must obtain between some in-progress P-eventuality and the result state ($\mathsf{TEL}(P)$) associated with P, but also that the process leading to $\mathsf{TEL}(P)$ must instantiate a complex network of internal causal relations, each of which must be verified in a context that instantiates a complete P-eventuality.

(iv) **The 'reverse' analysis** (adapted from Szabó, 2004, p.47). If P is telic and admits the progressive, then $\textsc{pfv}(P)$ is true iff:

a. $\textsc{prog}(P)$ is true of some event e
b. $\mathsf{TEL}(P)$ is true of some state s
c. e causes s
d. If e causes e' and e' causes s, then $\textsc{prog}(P)$ is true of e'

We are concerned with the analysis of telic progressives, which Szabó does not address directly. Even so, although we do not commit ourselves to the claim that the truth of a telic perfective is *constituted* by the actual causal relations in (iv), we take our account to be compatible with (iv), insofar as these conditions will follow from the truth of a telic perfective, at least for perfective operators that license culmination entailments.

[21]The philosophical literature employs the term *actual causation* to refer to token or singular instances of causation—i.e., an actual situation realizing both an effect and its cause(s)—and we use it in this sense throughout.

occur without requiring that anything takes place in the actual world: claims of actual causation are licensed by comparing actual events with relevant type-level information.

Crucially, then, if particular uses of accomplishments can license singular causal claims, these predicates must encode type-level information about the conditions under which their culmination conditions may be realized. This view is, moreover, compatible with the idea that certain components of a predicate (i.e., *processes*) may be instantiated as valid (partial) realizations of the predicate without requiring the culmination condition to occur. The upshot is that an uninflected telic predicate contains idiosyncratic (type-level) components relating both to culmination and to the processes that lead to it.

In what follows, we suggest that type-level causal information is naturally represented by a **causal model**. We begin with an overview of causal models and their relevance to singular causal claims, and then propose a formal treatment of telic predicates within the framework of *structural equation models* (SEMs; Pearl 2000).

3 Causal models for event types

3.1 Linguistic causation: the view from causal models

We suggested above that a *culmination procedure* (a type of *process*) should be understood in causal terms, and more specifically in terms of the kind of structure represented in causal models. We also highlighted one respect in which such models differ from token-level causal judgments, a point to which we return below. We now introduce a further clarification that is central to the present proposal. Causation in the world is not typically realized as a one-to-one relation between a single cause and its effect; rather, individual outcomes depend on the realization of multiple conditions. Our analysis will thus relate the notion of a *process* to the progressive realization of a series of conditions leading to an outcome, and a *culmination procedure* to a process whose outcome is given by the culmination condition of a telic predicate.

A multifaceted view of causation corresponds to cognitive reality. When we engage in practical, language-independent causal reasoning—for purposes of planning, intervention, explanation, and the like—we conceptualize outcomes as resulting from several prior factors. Whether a match ignites, for example, depends not only on its being struck on appropriate surface, but also on whether it is dry and whether oxygen is present. A realistic understanding of ignition includes the knowledge that, if the match is wet or oxygen is absent, no amount of striking will produce the desired effect. Causation thus appears as a many-to-one relation, in which an effect depends on a structured set of potentially interacting conditions. The term **causal model** can in principle refer to any structured representation of such dependency relations: after Section 3.3, we will use it to refer to *structural equation models* (Pearl, 2000).

The complexity of a cognitively plausible causal model stands in contrast with standard linguistic descriptions of causation, which express apparently binary relationships between single causes and their effects (see 23a-23c). This tendency aligns naturally with (and indeed, motivates) classical decompositional approaches (see, e.g., 2/3 in Section 1), on which causal expressions are uniformly analyzed as encoding a binary relation (such as CAUSE) between two events.[22] Once

[22]Decompositional approaches typically rely on dependency-based conceptions of causation, typically assuming a one-to-one relation between cause and effect. For a review of the literature, including discussion of the relevance of many-to-one approaches to causation for linguistic research, see Bar-Asher Siegal (2025).

linguistic descriptions are considered against the background of causal models, however, uniform CAUSE-decompositions face a number of difficulties. On the metaphysical side, reducing all causal dependence to a single relation immediately raises the familiar problem of *causal selection.* It is intuitively clear that (23a) provides a better causal description of a normal case of match ignition than (23b). Yet if a realistic conceptual model of match ignition treats both striking and the presence of oxygen as causally prior conditions that jointly contribute to the ignition event, it is far from obvious what principle should privilege the former over the latter (for further discussion, see Baglini and Bar-Asher Siegal, 2025, and references therein).

(23) a. Striking the match caused it to ignite.
b. ??The presence of oxygen caused the match to ignite.
c. The presence of oxygen enabled the match to ignite.

On the linguistic side, a uniform CAUSE-decomposition is difficult to reconcile with the diversity of causal expressions. Even granting an account of causal selection that explains the contrast between (23a) and (23b), how are we to explain the acceptability of (23c), if both *cause* and *enable* are assumed to introduce the same underlying causal relation?

If we assume that the practical concept of causation reflects cognitive reality, then causal models can be regarded as inputs to linguistic interpretation: i.e., objects that can be referred to, updated, and manipulated by standard linguistic means. Call this the **modeling perspective** on linguistic causation. On this view, descriptions of causal situations, like descriptions of any situation, need not encode the entire relevant structure. Rather, causal language picks out particular facets of a contextually salient causal model. The unifying feature of causal expressions is therefore not that they all lexicalize the same primitive causal relation, but that they all relate, in different ways, to a common kind of contextual structure. Thus, *cause* and *enable* in (23) may both be interpreted relative to the same model for match ignition, while differing in how they constrain the role played by their subjects within the set of conditions jointly responsible for the result. Similar ideas are developed by Nadathur and Lauer (2020) in their semantic analysis of periphrastic causative verbs, and by Baglini and Bar-Asher Siegal (2025) in their account of 'directness' inferences associated with lexical vs. productive causatives (see also Bar-Asher Siegal et al. 2021; Bar-Asher Siegal 2025).

The modeling perspective will prove explanatory only to the extent that it is paired with a representational framework capable of capturing the many-to-one structure of practical causal knowledge while also distinguishing among the different members of a causing set (see also Hobbs 2005): we employ structural equation models to provide linguistically relevant causal models. Developed by Pearl (2000) and subsequent work (Spirtes et al., 1993; Steyvers et al., 2003; Halpern and Pearl, 2005; Halpern, 2015, a.o.), SEMs have been used extensively in the study of causal inference, explanation, and decision-making across the cognitive sciences, and are therefore strong candidates for the formal representation of practical causal knowledge.[23]

3.2 Causal claims and causal models: token vs. type causality

The causal modeling perspective on linguistic causation regards causal models as repositories of causal information that may be relevant for non-linguistic purposes (e.g., deciding how to act in a

[23] Our choice of framework is also motivated by the recent use of SEMs to analyze linguistic phenomena including *counterfactual* conditionals (Schulz 2011; Henderson 2014; Bjorndahl and Snider 2015; Ciardelli et al. 2018), *causative* claims (see references above), and expressions not obviously causal in nature, such as *implicative* and *abilitative* predicates (Baglini and Francez 2016; Alonso-Ovalle and Hsieh 2021; Copley 2021; Nadathur 2023a,b).

given situation) as well as linguistic ones (e.g., deciding how to describe the situation). A model is a language-independent cognitive object encoding a set of generalizations that reflect what is known or believed about how causality works in a given system. These generalizations (model-internal relationships) may produce causally-driven expectations in particular contexts, but they are not *about* any particular context.

By contrast, episodic causal claims like (23)-(24) explicitly focus on singular contexts: such statements adhere to the general pattern of non-modal declaratives, describing concrete scenarios (i.e., *what happened*) at specific spatiotemporal locations. The apparent truth conditions of a periphrastic causative claim are schematized below: (24) requires the realization of both a causing (vase-dropping) eventuality and a result (vase-breaking), and also that the former led to the latter.

(24) Ria dropping the vase caused it to break.

a. Ria dropped the vase. (event c)

b. The vase broke. (event e)

c. Event c (is part of what) *brought about* event e. (dependency condition)

The difficulty with claims of singular causation lies in dependency conditions like (24c) (cf. Hume 1748). While (24a)-(24b) report on concrete, directly observable facts, (24c) does not. We cannot, however, treat a condition like (24c) as irrelevant for (24), since it is intuitively clear that "unlinked" occurrences of c and e are insufficient for the truth of (24), even if both events take place within reference time and in the correct temporal order ($e \not\prec c$). To the extent that semantic interpretation relies on the verification of particular atomic predications, the truth conditions in (24) necessitate the existence of a reference object (a *witness*) for the causal relationship in (24c): this object must be something which goes beyond observables at the referential here-and-now. From the causal modeling perspective, this is precisely what a (salient, contextually-relevant) causal model supplies: an abstract but structured object against which causal predications can be checked.[24]

The idea is that causal models provide licensing conditions for causal claims: the truth and appropriateness of (24) depends not only on the occurrence of events c and e, but also on the availability of a model which encodes a causal relationship (i.e., a generalization) linking occurrences of events of the c type to events of the e type.[25] On the modeling perspective, a suitable formal framework for describing the structure of a causal model should allow such a model to act as a bridge between world knowledge (of a causal nature) and the familiar objects of formal semantics, representing causal knowledge in a format which allows us to express precise (set- and/or model-theoretic) conditions under which a relation corresponding to (24c) is verified.

Causal models are thus not "about" causal statements, but reflect what individuals know about the world. Although we hold that singular causal claims like (24) rely on causal models, we also do not claim that these statements are "about" causal models. Singular causal claims are, as they seem to be, about particular situations in the world: they are simply made *in view* of situationally-relevant causal knowledge. The distinction between causal models and singular claims can be

[24] This is not the only possible response to the problem posed by causal 'unobservability'. There are various alternative reductionist approaches, including counterfactual approaches, on which the causal relation in (24) is decomposed into (alethic) modal statements (see, e.g., Lewis, 1973a; Sosa and Tooley, 1993). To the extent that such approaches cohere with CAUSE-decompositions, they are often presupposed in linguistic work (e.g., Dowty, 1979). Bar-Asher Siegal and Boneh (2020) discuss the relevant philosophical literature and its points of connection and contention with the linguistic literature and data; see also Copley and Wolff (2014).

[25] Following Hume (1748), Davidson (1967) similarly treats causation or causal laws as a set of *nomological generalizations*: i.e., statements of lawful regularity that may be inferred through experience.

likened to the distinction between **type-** and **token-level statements**. Type-level claims are a type of generic statement, expressing overarching patterns or 'laws' in the world: causal generics like "Smoking causes cancer," which seem to express relations between categories of events, illustrate this in the causal domain. Token claims are necessarily episodic, describing individual occurrences at specific places and times. Singular (e.g., past-tense) causal claims are of this sort, emphasizing a local cause-effect link but failing to describe general links between categories of events.[26]

From a scientific (empirically-driven) point of view, of course, singular causal claims describe the sorts of real-world situations which must be regularly observed in order to formulate causal generalizations, and thus to license acceptable type-level causal claims. The flow of comprehension, then—i.e., the development of causal knowledge—must run from the specific to the general, or from token instances of causation to type-level models: from this perspective, a model is simply a reification of causal regularities drawn from robust observation of specific causal instances (Hausman, 1998, 2005). A causal model is thus constructed on the basis of observation and experimentation; the statistical and qualitative correlates of model construction are a key research area in cognitive science (see Woodward 2003; Hitchcock 2020), but will not concern us here.

It is the reverse perspective—moving from the general to the specific—which is our focus. Once a causal model is established (has become cognitively available), the causal generalizations which it encodes correspond to a set of *nomological* relations between certain types of objects (conditions, properties, punctual events), by virtue of which token-causal claims may be verified.[27] The truth conditions of an individual instance of causation—specifically, a condition analogous to (24c) for (24)—are formulated in terms of these type-causal relations (see also Tooley, 1987; Hoover, 2001). As a result, even though the causal relation itself is not observable, causality can be attributed to specific states of affairs based on an understanding of type-level relations. For instance, (24) reports on an actual vase-dropping action by Ria, a state transition involving the vase, and further indicates that the occurrences of these events conform, in their relative spatiotemporal locations, to a nomological relation between events of the dropping- and state-transition types in a causal model. In order for her to felicitously utter (24), this causal model must be (a) available to the speaker and (b) presumed by her to be shared with her audience (cf. Baglini and Bar-Asher Siegal 2025; Bar-Asher Siegal 2025). This causal model encodes knowledge about event types—specifically, general principles governing the culmination of processes. As noted above, type-level knowledge is crucial for evaluating, at an intermediate stage of a process, whether it is on course to reach its intended culmination. This idea will play a central role in our semantic analysis of telic progressives.

3.3 The formal framework: structural equation models for causation

Practical causal reasoning requires three things: an actual context or situation, a set of salient or relevant causal laws (a model), and a mechanism for applying the latter to the former. We use the Lifschitz circuit context in (37a) to illustrate: (25a) states the relevant causal laws, and (25b) provides a context in which these laws are presumed to be in effect.

(25) **The Lifschitz circuit.** Suppose we have a circuit with two switches and one light.

[26] See also Gehrke and McNally (2009) for related discussion of the link between event types and verb phrases.

[27] See also McHugh (2023), who argues that causal models must represent nomic possibility, in the sense that they determine, for each possible world, whether it is nomically (im)possible. In our approach, this is not a further requirement on causal models, but part of what constitutes them as causal models.

a. The light turns on just in case both switches are set to the same position: either both up or both down.

b. At the moment, switch 1 is down, switch 2 is up, and the light is off.

Given (25a), (25b) represents a *causally normal* state of affairs: the state of the dependent element (the light) is as expected given the states of the independent elements (the two switches). The information in (25a) can also be used to make predictions: namely, that flipping either of the switches will produce a change in the light. The availability of such predictions also allows us to establish links between flipping the switches and turning the light on or off, and to make specific causal claims about the relations between these events.

The information in (25a) expresses the type-level causal relations which should be represented in a model for a Lifschitz circuit. Figure 1 represents these relations more formally. The graph in Figure 1a conveys that (the state of) the light L causally depends on (the states of) the switches S_1 and S_2. The three-valued truth table in Figure 1b specifies the nature of the dependencies, indicating that the light comes on (L takes on value 1) just in case S_1 and S_2 have the same value (both 0 or 1; let 0 stand for 'off' and 1 for 'on'). The third truth value, u, represents that a node's status is unknown; as shown in the truth table, an indeterminate (u) value for either switch results in an unknown value for the light.[28] The causal normality of the situation in (25b) is reflected in a correspondence between the actual values for S_1, S_2, and L (0,1, and 0, respectively) and the sixth line of Figure 1b; the causal predictions mentioned above can similarly be read from 1b by moving to a line in which one or the other of the switches changes its value, and noting that the corresponding value of L changes in each case.

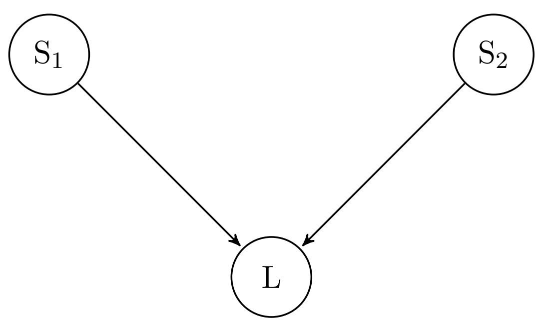

(a)

$f_L =$

S_1	S_2	L
u	u	u
u	0	u
u	1	u
0	u	u
0	0	1
0	1	0
1	u	u
1	0	0
1	1	1

(b)

Figure 1: Graphical model for the Lifschitz circuit

Figure 1 is an example of a **structural equation model**. The skeleton of an arbitrary SEM comprises a *directed acyclic graph* $D = \langle V, A \rangle$ in which the set of vertices V is given by a finite set of propositional variables, and the set of directed edges corresponds to a relation $A \subset V \times V$ of ordered pairs of vertices, corresponding in this context to an atomic notion of *causal influence*. Vertices can be paired with *truth values* from the set $\{u, 0, 1\}$; a value of 0 or 1 identifies a propositional variable

[28]The use of a three-valued epistemic system is consistent with our aim of modeling speakers' causal knowledge within the SEM framework. Epistemically, causal knowledge concerns the inferences that speakers are able to draw about causal relations, even in the absence of complete information about all relevant conditions.

with a *proposition* in the familiar linguistic sense, while u indicates that the value of the vertex is *undetermined* (epistemically unsettled). Given two vertices $X, Y \in V$, $\langle X, Y \rangle \in A$ indicates that X is an immediate causal ancestor of Y: more precisely, that the value of X influences the value of Y. We distinguish a set $I \subseteq V$ of *inner* or *dependent* variables: $I = \{X \in V \mid \exists Y\, [\langle Y, X \rangle \in A]\,\}$ (see also Schulz, 2011). In the model in Figure 1, the set V is given by $\{S_1, S_2, L\}$, the relation A by $\{\langle S_1, L \rangle, \langle S_2, L \rangle\}$, and the set I by $\{L\}$.

The graph shows us where causal dependencies exist, but does not establish their nature. This skeleton must therefore be paired with a specification of *how* the values of dependent variables are influenced by the values of their ancestors. We let this be handled by a (partial) function F which maps each dependent vertex $X \in I$ to a pair $\langle Z_X, f_X \rangle$ where $Z_X = \{Z : Z \in V \wedge \langle Z, X \rangle \in A\}$ is the set of X's *immediate causal ancestors* and $f_X : \{u, 0, 1\}^{|Z_X|} \rightarrow \{u, 0, 1\}$ is a function from valuations of Z_X to valuations of X.[29] As shown in Figure 1b, the function f_X for a given dependent variable X can be expressed in the form of a truth table. It might, equivalently, be given as the (one-directional) logical form (or *structural equation*) in (26); (26) follows the strong Kleene logic for conjunction, as per Figure 1b.[30] Given a *situation* s—a three-way assignment of truth values to V (see Definition 28b)—we can use the information in F to check for causal consistency (adherence to the model) as well as to predict the causal consequences of any changes to s.

(26) Structural equation (f_L) for the Lifschitz example: $L := S_1 \wedge S_2$

The structure of an SEM allows us to define a variety of model-internal relationships. From the causal modeling perspective, the potential referents of causal expressions are well-defined types of structural dependency which may obtain between an inner variable X and one or more of its causal ancestors (informally, variables that are 'upstream' of X in the relevant model).

(27) **Causal ancestors.** Let $M_D = \langle D, F_D \rangle$ be a causal model over a set V of propositional variables. The set $\mathsf{Anc}(X)$ of *causal ancestors* for any $X \in V$ is given by $\{Y \in V | R_A(Y, X)\}$, where the relation R_A is the transitive closure of A (the set of directed edges of D).

We are interested in formalizing the notion of a **culmination procedure** or 'recipe' for a given result: intuitively, this should constitute a causal pathway which picks out a set of *facts* (Definition 28a) that work together (i.e., are collectively *sufficient*) to produce some target effect, and which are independently required (*necessary*) for this effect (when considered as part of the pathway). The definitions in (28) provide a basis for formalizing these intuitions by allowing us to refer to variable-value pairs (facts), particular value assignments (situations), and the variables which are assigned 0 or 1 truth values in these assignments. A *fact* corresponds to the familiar linguistic notion of a proposition; note that a *situation* (Definition 28b) is uniquely identified by its *kernel* (the set of facts it establishes; Definition 28e), but not by its *domain* (Definition 28c).

(28) Let $M_D = \langle D, F_D \rangle$ be a causal model over a set V of propositional variables.
   a. **Fact.** A *fact* is a pair $\langle X, x \rangle$ where $X \in V$ and $x \in \{0, 1\}$.
   b. **Situation.** A *situation* $s : V \rightarrow \{0, 1, u\}$ is an assignment of truth values to V.

[29] Following Schulz (2011), linguistic applications of SEM typically assume that F should be *rooted* in the set $B = V - I$ of *exogenous* (or *outer*) variables: this ensures that any path which traces backwards through the dependencies of a given variable will terminate in B, and prevents circular chains of causation. We will assume rootedness here; see Schulz for a more formal treatment of this property.

[30] This differentiates our framework from that of Schulz, who uses a weak Kleene logic for causal consequences.

c. **Domain.** The *domain* $\mathsf{Dom}(s)$ of a situation s is the set of propositional variables to which s assigns a value other than u: $\mathsf{Dom}(s) = \{X \in V | \langle X, 1\rangle \in s \vee \langle X, 0\rangle \in s\}$.

d. **Determination.** If $X \in \mathsf{Dom}(s)$, we say that X is *determined* by s.

e. **Kernel.** The set of determinations made by a situation s is called its *kernel*: $\mathsf{Ker}(s) = \{\langle X, s(X)\rangle | X \in \mathsf{Dom}(s)\}$

f. **Supersituation.** Given two situations, s and s', s' is a *supersituation* of situation s just in case $\mathsf{Ker}(s') \supseteq \mathsf{Ker}(s)$.

Model structure induces a notion of *causal consistency* (Definition 29). A situation s is consistent with respect to model M just in case it is causally normal relative to M—i.e., if the values it assigns to a dependent variable X match the predictions of f_X, given the values s assigns to $\mathsf{Anc}(X)$. Thus, the situation $s = \{\langle S_1, 0\rangle, \langle S_2, 1\rangle, \langle L, 0\rangle\}$ in (25b) is causally consistent with respect to the Lifschitz model in Figure 1, while the alternative situation $s' = \{\langle S_1, 0\rangle, \langle S_2, 1\rangle, \langle L, 1\rangle\}$ is not (since it fails to conform to the structural equation f_L in Figure 1b).

(29) **Causal consistency.** Let $M = \langle D, F_D\rangle$ be a causal model over a set V. Situation s is *causally consistent* iff $\forall X \in \mathsf{Dom}(s) \cap I(\subset V)$, we have $s(X) = f_X(\overrightarrow{s(Z_X)})$

The notation $\overrightarrow{s(Z_X)}$ used above refers to the sequence of valuations assigned in situation s to the set Z_X of ancestors of variable X; this provides an appropriate input to the function f_X.

Causal consistency feeds the definitions of **causal necessity** and **sufficiency** in 30 (see also Baglini and Francez 2016; Nadathur and Lauer 2020; Nadathur 2023b; Baglini and Bar-Asher Siegal 2025). These relations hold only between (collections of) facts, and thus can only be assessed in the context of a truth value assignment or *situation*; moreover, since it must be the causal information in an SEM that determines whether necessity and sufficiency obtain, these definitions must make reference to (hypothetical) situations that conform to the model.

(30) a. **Causal necessity** (with respect to a situation). Let $M_D = \langle D, F_D\rangle$ be a causal model over a set V of propositional variables, and let s be a causally consistent situation. Fact $\langle X, x\rangle \in \mathsf{Ker}(s)$ is **causally necessary** in s for fact $\langle Y, y\rangle \in \mathsf{Ker}(s)$ iff $X \in \mathsf{Anc}(Y)$ and for any causally consistent situation s' such that $\mathsf{Dom}(s) = \mathsf{Dom}(s')$, and $s(X) \neq s'(X)$, we have $s - s' \supseteq \{\langle X, x\rangle, \langle Y, y\rangle\}$. We write $\langle X, x\rangle \lhd_s \langle Y, y\rangle$.

b. **Causal sufficiency.** Let $M_D = \langle D, F_D\rangle$ be a causal model over a set V of propositional variables, and let s be a causally consistent situation. Situation s is **causally sufficient** for $\langle Y, y\rangle \in \mathsf{Ker}(s)$ iff for all $X \neq Y \in \mathsf{Dom}(s)$, $\langle X, s(X)\rangle \lhd_s \langle Y, y\rangle$. The set $S = \mathsf{Ker}(s)$ is a **sufficient set** for $\langle Y, y\rangle$: we write $\mathsf{Suff}_{M_D}(S, \langle Y, y\rangle)$.

Per Definition 30a, *causal necessity* holds between two facts in a consistent situation just in case changing the value of the 'upstream' fact would require a corresponding change in the downstream fact in order to maintain consistency: in other words, toggling the upstream node forces changes to the downstream one. *Causal sufficiency* (Definition 30b) relates a consistent situation to a component fact: s is causally sufficient for $\langle Y, y\rangle \in s$ if each of the (other) facts in s is necessary for $\langle Y, y\rangle$ with respect to s and these facts collectively predict the value that s actually assigns to Y when taken together with the relevant model.[31]

[31] Definition 30b embeds a notion of *minimality* (cf. Baglini and Bar-Asher Siegal 2025), in excluding from the set of sufficient situations for $\langle Y, y\rangle$ any situation whose domain includes variables of which Y is causally independent (i.e., any $X \in V$ such that $X \notin \mathsf{Anc}(Y)$). This follows from the necessity constraint in 30b. While it is possible for

As defined by a model M, a sufficient situation s for fact E can be construed as a potential 'recipe' or procedure for the (actual) realization of E. It need not, however, be the *only* such procedure specified by M. The model in Figure 1, for instance, provides two (disjoint) sufficient sets for the light to be on: $\{\langle S_1, 0\rangle, \langle S_2, 0\rangle, \langle L, 1\rangle\}$ and $\{\langle S_1, 1\rangle, \langle S_2, 1\rangle, \langle L, 1\rangle\}$ (lines 5 and 9 of Figure 1b, respectively). Practically speaking, the light is expected to be on as long as one of these sets is actualized, making each set individually sufficient but unnecessary for the designated effect. The model also defines two sufficient sets for the light to be off: $\{\langle S_1, 0\rangle, \langle S_2, 1\rangle, \langle L, 0\rangle\}$ and $\{\langle S_1, 1\rangle, \langle S_2, 0\rangle, \langle L, 0\rangle\}$ (lines 6 and 8 of Figure 1b, respectively).

Within a sufficient set for effect E, necessary conditions are defined by the role they play with respect to E; as such, they can be understood as sites of potential intervention in an attempt to bring about (or prevent) E's realization. The scenario in (25b) describes one of two sufficient sets for the Lifschitz light to be off ($\langle L, 0\rangle$); the determination $\langle S_1, 0\rangle$ for S_1 is necessary for this effect within the actualized sufficient set, since any change to the value assigned to S_1 would require a change to the value of L in order to maintain consistency. The fact $\langle S_1, 0\rangle$ is thus necessary for the light to be off within a sufficient set of conditions for this result. This makes it a potential cause for the actualized situation in (25b), but (as per the discussion in Section 3.1) the truth of a claim like "Flipping switch 1 off turned the light off" would also require the actualization of both a switch-flipping and light-changing event (with the latter occurring no earlier than the former).

Definitions (30a)-(30b) echo Mackie's (1965) treatment of causality, albeit in a non-alethic setting (see Baglini and Bar-Asher Siegal 2025 for discussion). Mackie suggests that singular (single-event) *causes* can be reconstructed as so-called 'INUS' conditions, comprising Insufficient but Necessary elements of an Unnecessary but Sufficient set for their results: (30a)-(30b) give content to the INUS concept in the SEM framework. Defining necessity and sufficiency in model-theoretic terms grounds these relations in information that can be extracted from a causal model, without reference to a particular context of evaluation: they express **type-level causal** relationships, corresponding on the one hand to knowledge of how certain valuations for dependent variables in the graph can be realized, and on the other to potential truth-makers for singular causal claims in any actual context involving variables of the type included in the model.

3.4 Connecting causal models to telic predicates

On our implementation of the modeling perspective, both sufficient sets and their component elements—that is, individual necessary conditions—may serve as legitimate anchors for linguistic descriptions of actual causation. An episodic causal claim ϕ, linking a causing condition C to a result E, requires not only that C and E be realized, but also presupposes the existence and cognitive availability of a model in which C is a member of some sufficient set for E.

Returning to the relation between causal claims and causal models discussed in Section 3.1, commitment to the truth of ϕ further commits the speaker to the claim that E was realized *via* the actualization of a sufficient set containing C. It follows that ϕ is infelicitous in discourse contexts incompatible with the realization of the full sufficient set of which C is a necessary member. Distinct causal expressions plausibly impose distinct additional constraints on the role that C plays within an active sufficient set for E.[32] On the present view, however, such expressions are unified in

a 'larger' situation s' (that is, a situation s' such that $\mathsf{Ker}(s') \not\subseteq \mathsf{Anc}(Y)$) to be causally relevant for $\langle Y, y\rangle$ in a given discourse context, this relevance can only hold where s' is a supersituation of a causally sufficient situation as defined above; Definition 30b can therefore be adopted with no meaningful loss of generality.

[32]Thus, so-called direct causatives may require C to complete a sufficient set, whereas indirect causatives do not

crucial respects: they all make reference to a model in which C and E are linked by the relations in Definition 30. The truth and felicity of singular causal claims thus rely on causal models—that is, on representations of type-level causal information—as proposed in Section 3.1.

We suggest that telic predicates introduce just this sort of type-level causal information. One piece of evidence for this proposal was sketched in Section 2.4: when an accomplishment predicate P is used in a culminating context (e.g., with the English simple past), it commits the speaker not only to the realization of the culmination condition C_P, but also to the occurrence of a process that led to C_P, that is, to the realization of some set of conditions jointly sufficient for C_P. From (31), we infer not only the truth of the culmination condition in (31b), but also the existence of a prior progressive instantiation of the underlying predicate (see 31a), and a causal dependency (see 31c), between that instantiation and the relevant culmination; compare (31) with (24).

(31) Emanuel baked a cake.

a. Emanuel engaged in some activity, or series of activities. (c) (PROG(P))

b. A cake came into being. (e)

c. c is (part of) what brought about e.

On our analysis, verifying a relation like (31c) requires reference to a causal model. This implies that culminated uses of telic predicates depend on type-level causal information about the conditions under which C_P is realized. In the case of singular accomplishment claims, just as in ordinary causative claims such as (24), verification of an actual causing event—here, the progressive-satisfying event c—requires appealing to type-level information. The event c must instantiate (part of) a fully realized sufficient set for C_P, and the verification conditions for PROG must, as argued in Section 2.4, target procedural information concerning how events of type C_P are brought about. Our claim is that this procedural information is introduced by the accomplishment predicate itself: used felicitously, any predicate in this aspectual class makes reference to a speaker's normative world knowledge concerning how events of the relevant culmination type are typically realized.

We therefore propose that an accomplishment *event type* corresponds to a causal model: that is, to a body of generalizations about the complex causal interrelations among conditions—events, states, properties, facts, and the like—that jointly determine realistic 'recipes' for C_P. An uninflected telic predicate P introduces a causal model M_P in which P's culmination condition, C_P, appears as a dependent variable. Model structure yields a natural interpretation of a *culmination procedure*: the culmination procedures associated with P are simply the causal pathways (i.e., the sufficient sets) for the truth of C_P as encoded in M_P. Depending on the predicate, the model may specify a single culmination procedure S ($\mathsf{Suff}_{M_P}(S, \langle C_P, 1\rangle)$) or several distinct procedures (see Section 4.2). Verification of a culminated instance of P, as in (31), then requires only that one such sufficient set be fully realized, thereby validating (31a–31c).

As noted in the introduction, our focus in this paper is on developing a semantics for a progressive operator in the causal modeling perspective, as a means of demonstrating the viability (and the advantages) of such an approach. We do not offer a full account of how a given predicate comes to be associated with a causal model, leaving this as a matter for future work. Some preliminary comments about this process—and about the sense in which a relevant model is "invoked" by the use of a predicate—are nevertheless in order.

First, we do not suggest that a predicate can only be associated with a single model, even by an individual speaker. We have in mind instead a system in which certain features of an appropriate

(see, e.g., Martin 2018; Nadathur and Lauer 2020; Bar-Asher Siegal et al. 2021; Baglini and Bar-Asher Siegal 2025).

model are fixed by the lexical components of a predicate, while others are influenced by discursive or even non-linguistic context. For instance, a model associated with the predicate *open a/the door* necessarily involves a door which undergoes a state transition, as well as an agent or effector that engages with a particular mechanism linked to the door. However, whether this mechanism involves a doorknob, an electronic “door open” button, a keypad, or even some combination of the above—and thus the set and specifics of the culmination pathways that the predicate model provides—will depend on the context in which the predicate is used, and in particular on any features of the door in question that are contextually apparent to a speaker. It may thus be appropriate to say, as suggested by a reviewer, that a given predicate is lexically associated (linked in the mental lexicon) with a *class* of causal models, rather than a singular model. Nevertheless, we believe that it is indeed a singular (albeit type-level) model that is invoked and composed with the progressive (or any other aspectual operator) in a given use context, and that this is precisely that model available to the speaker that best suits their understanding of the situation under description.[33]

If specifics of an utterance context can be involved in the construction of a model, it is reasonable to ask in what sense the relevant model constitutes a *type*. The important point here is that the specifics in question—even when they take into account individualized features of a given situation—appear to be systematically restricted. In particular, situational specifics relevant for model construction are just those that affect the nature of the task and routes to task completion: incorporating this information nevertheless allows the model to define a set of qualifying events.

Although there is much investigation to be done, we believe that this view gains empirical support from judgments of telic progressives which suggest that specifics of a predicate’s internal argument are taken into account by the event type model, but that those of the external argument (beyond selectional restrictions imposed by the verb) are not. This can be illustrated by comparing the example in (32) (loosely adapted from Asher 1992) to the case of Benny, our undertrained would-be ultramarathoner from example (19).

(32) **The impassable barrier scenario.** Suppose that we are looking at a street which is bisected lengthwise by an extremely tall barrier which is known to be highly electrified. Henrietta begins to walk towards the barrier from one side of this street.

a. #Henrietta is crossing the street.

In (32), specific knowledge of the street referenced in the verbal predicate renders the target claim (32a) infelicitous, on par with the IE progressives discussed in Section 2. The reason for this is intuitively clear: although a street with an impassable barrier still counts as a street, its specific properties render the street-crossing task unachievable. It is thus infelicitous to say that Henrietta is crossing this street for the same reason that it is infelicitous to say that any human is swimming across the Atlantic Ocean—there is simply no way for a human to complete either task. Just as task impossibility is established via knowledge about the Atlantic in the swimming example, impossibility in (32) is established via specific knowledge of the street (i.e., the internal argument). If instead of an impassable barrier, we imagine a street with a low bisecting fence, the crossing task once again becomes possible, and this is because the event type model invoked in assessing (32a) is now one that reflects the presence of this fence: in this scenario, the model will establish that sufficient sets for culmination involve some means of bypassing the fence (by climbing

[33] In the absence of any specifics, it may be that a “coarse” model is invoked—i.e., one which has only the features that are shared by the class of appropriate models—but we suspect instead that speakers and hearers select what they take to be an exemplar or canonical model. In the case of an arbitrary door-opening, this is likely to be a mechanical rather than automatic door.

over it, using a gate, etc.). The actual properties of the street correspond to information that must be included in the model, since it is only in view of matching reference time activities to the nature of this particular street-crossing task that a progressive claim can be truth-conditionally evaluated.

Scenario (32) contrasts with ORC cases like Benny's, which involve in-principle possible tasks that are simply locally unachievable for a specific protagonist. The empirical judgment in (19a)—in contrast with (32a)—illustrates clearly that Benny's individual properties do not automatically render the progressive infelicitous: on our analysis, this is because Benny's actions at reference time can still be compared to a well-defined event type model for ultramarathon running (which may indeed take into account specifics of the course which he undertakes). The key point is that while Benny's training determines whether or not it is possible *for him* to complete the task in a given context, his local properties do not affect the nature (human possibility and/or the specific set of culmination pathways) of the task itself. This distinction—between relevant situational specifics (e.g., of the internal argument) that determine the nature of the task, and specifics (e.g., of the external argument) which may affect the outcome of particular (singular) courses of events—allows us to talk about individual models invoked by telic predicates in given utterance contexts while simultaneously treating these models as repositories of (event) type-level information.[34]

This brings us back to the starting point of the paper. On the decompositional approach to accomplishment predicates represented in (2) (see Section 1), the idiosyncratic content of a predicate resides exclusively in its result state. The preparatory stages are left un(der)specified, encoded only by a predicate such as ACT. On the present analysis, by contrast, the denotation of an accomplishment must also include idiosyncratic information about the process that leads to culmination, namely, information that determines what counts as a normative pathway to culmination.[35]

3.5 Conceptualizing *process*

The central difficulty in stating truth conditions for telic progressives is that such conditions appear to depend simultaneously on intensional information (in view of their connection to culmination) and on a partitive notion grounded in actual facts. The causal modeling perspective allows us to reconcile these two demands. The type-level causal information encoded in a model is inherently intensional, with reference to culmination supplied by culmination procedures in M_P. At the same time, culmination procedures correspond to sets of facts (the kernel of a sufficient situation for culmination in M_P) which can be mapped to spatiotemporally located *situations* (i.e., situations in the familiar possible-worlds sense) by virtue of the propositions which these sets contain. This makes it possible to compare reference time facts with culmination procedures in set-theoretic terms. If the reference time facts realize part, but not all, of a culmination procedure for C_P, they may constitute a *partial realization* of P.

The tension between existing approaches to the imperfective paradox can be understood as

[34] The observation that properties of the internal but not external argument affect event type structure aligns with existing work on syntactic constituency (see Kratzer 1996 and citing work); we suggest that this asymmetry not only motivates the notion of a singular albeit type-level model but may also constitute a useful starting point for investigating the link between model construction and predicate-internal composition.

[35] The term *P-eventuality* standardly refers to an event in the denotation of P. We retain the term here, but reinterpret it within the present framework. Since the denotation of a telic predicate P corresponds to a causal model for C_P, we use *P-eventuality* to refer to complete or partial realizations of culmination pathways for C_P. Such realizations will often be complex, involving multiple interacting conditions, and may themselves contain substages or subeventualities, some but not all of which may independently satisfy P. In this sense, the causal-modeling approach induces a causal mereological structure internal to the predicate, which then serves as input to aspectual operators.

reflecting different answers to the following question: how should the *process* component of a complex telic event be characterized? The literature distinguishes two main strategies.

(A) **Ontological approach.** Following Stout (1997), this approach treats processes and events as distinct ontological categories. A process, picked out by a nominalization such as *Emanuel's cake-baking*, excludes the completed outcome, whereas a completed event or culmination, as picked out by *Emanuel's baking of the cake*, is a bounded event. On this view, the progressive directly denotes a process entity, while the perfective denotes a culmination-bearing event.

(B) **Mereological approach.** Following Landman (1992) and Bonomi (1997), a process is analyzed in mereological terms, as a proper part or series of parts of a larger culminating event. The question is then which subevents genuinely count as parts of a larger eventuality denoted by the predicate.

Our view is closer in spirit to the mereological approach. However, once the denotation of a telic predicate is reconceived as a type-level causal model, there is no need to explicitly formulate the notion of *process* as a mereological relation between events. In our framework, a process is neither a *sui generis* ontological category nor merely a proper part of a larger event. Rather, it is a *causally structured assembly of states and events* that aligns with one of the sufficient sets for culmination. Each element in such an assembly is necessary relative to at least one causal pathway for C_P, and the elements of a pathway are jointly sufficient for realizing C_P under normal conditions.

It follows that processes may be instantiated, or fail to be instantiated, in particular contexts. As we spell out in Section 4.1, a reference time situation verifies the progressive with respect to P (i.e., verifies that P is *in progress*) just in case it instantiates a proper part of a process for C_P by satisfying some, but not all, of the conditions in a causally sufficient set, while not satisfying a sufficient set for non-culmination. The truth of a telic progressive thus requires more than partial realization: it must also not be the case that reference time facts falsify or causally preclude C_P. Crucially, however, this further requirement can also be stated model-theoretically: as noted in Section 3.3, the structural equations for C_P and its causal ancestors in M_P determine not only sufficient situations (culmination procedures) for $\langle C_P, 1\rangle$, but also sufficient sets for $\langle C_P, 0\rangle$ (non-culmination procedures). Consequently, although verification of a telic progressive requires reference to intensional objects defined by a causal model, these objects can be compared *as sets* to reference time facts, making good on the core partitive intuition.

The causal model concept of *process* makes it possible to characterize processes both causally and compositionally, without appeal either to projected culmination or an intensional analysis of the progressive aspect. By replacing the notion of a mereological relation between events with causal model structure, we integrate partitivity and causal adequacy into a single conception of process. Before turning to the formal truth conditions for telic progressives, it will be useful to restate the role of type-level causal information in our account:

1. A telic eventuality of type P is *in progress* just in case some culmination procedure is partially realized at reference time, while no non-culmination procedure is fully realized. Relative to model M_P, an in-progress P-eventuality therefore necessarily instantiates a *cause* of C_P: namely, some condition or set thereof belonging to a sufficient set for C_P.

2. The actualization of an in-progress P-eventuality neither predicts nor requires the realization of C_P itself, precisely because its relation to C_P is defined only at the type level (in model M_P) and is not constituted by an actual causal relation to culmination.

This is how the causal modeling perspective resolves the imperfective paradox. Because type-level causal relations are generalizations, actual situations may instantiate a normative process for C_P without actually bringing C_P about, making it possible to truthfull and felicitously use a telic progressive in a non-culminating context.

4 Towards a semantics for telic progressives

4.1 Truth conditions for telic progressives

Truth and felicity conditions for telic progressives can now be stated by using the machinery from Section 3 to formalize the descriptive generalization (restated in 33) from Section 2.4. This generalization captures what it means for a telic eventuality to be *in progress*: in light of the previous discussion, this requires the partial but non-terminal realization of a culmination procedure.

(33) **Descriptive generalization for telic progressives.** Let P be an accomplishment predicate with culmination condition C_P. Then:

a. $\textsc{prog}(P)$ is felicitous iff P is associated with at least one culmination procedure.

b. Where felicitous, $\textsc{prog}(P)$ is true iff the reference time situation corresponds to a strictly partial realization of a culmination procedure, reflects a change in this procedure as compared to the pre-reference time situation, and does not settle that C_P is false.

Section 3.3 defines a culmination procedure for telic P in terms of a sufficient set for the truth of culmination condition C_P in a model M_P in which C_P is a dependent variable. (33a) thus amounts to the requirement that P denote a non-empty causal model: i.e., that there is a structured causal representation for events of this type. Given a model, a P-eventuality is in progress at time t just in case the conditions stated formally in (34) hold. First, some culmination procedure for C_P must have been initiated at t, reflecting a reference-time change: the reference situation must verify at least one fact in a sufficient set for culmination which did not hold immediately prior to reference time. Second, no culmination procedure has been completed at t: the reference time situation leaves unsatisfied at least one condition in each sufficient set for culmination ($\langle C_P, 1\rangle$). Finally, no sufficient set for the falsity of C_P has been completed at t: the reference situation leaves unsatisfied at least one condition in each sufficient set for non-culmination ($\langle C_P, 0\rangle$). Condition (i) guarantees instantiation of a (non-maximal) P-eventuality; (ii) and (iii) jointly ensure that this eventuality has neither culminated nor terminated some other way, so that it is *ongoing* at reference time. In (34), s and S represent situations (sets of *facts*); Q picks out individual facts (see Definition 28a).

(34) **Analytical truth conditions for telic progressives.** Let M_P be a model for telic predicate P with culmination condition C_P. Given a time t and a world w, let w_t be the set of facts (true propositions) in w at time t, and $s(M_P, w, t)$ be the assignment of values to nodes of M_P that corresponds to what is known in w at time t ($\mathsf{Ker}(s(M_P, w, t)) \subseteq w_t$). $\textsc{prog}(P)$ is true at w, t just in case the following conditions are satisfied:[36]

[36] We assume a standard treatment of times as densely ordered and treat *situations* and *worlds* as assignments of values to propositional variables, in which only a situation can permit the value u (see Section 3.3). Although this differs from the conception of situations and worlds established in a possible worlds framework, the two notions are compatible, insofar as a situation in our sense establishes 0-1 values for a subset of the propositions determined by a world and can thus be regarded as *part* of one or more possible worlds at specific times; this relationship is captured by writing $\mathsf{Ker}(s) \subseteq w_t$ for time t. Situations as defined via causal models are not inherently time-dependent, but since they are made up of propositions, they can hold or fail relative to particular times in an evaluation world.

(i) $\exists S\,[\mathsf{Suff}_{M_P}(S, \langle C_P, 1\rangle) \wedge$
$[\exists Q\,[Q \in (S \cap s(M_P, w, t)) \wedge \exists t' \prec t\,[\forall t''\,[t'' \in [t', t) \rightarrow Q \notin s(M_P, w, t'')]]]]]]$

(ii) $\forall S\,[\mathsf{Suff}_{M_P}(S, \langle C_P, 1\rangle) \rightarrow [\exists Q\,[Q \in (S - s(M_P, w, t))]]]$

(iii) $\forall S\,[\mathsf{Suff}_{M_P}(S, \langle C_P, 0\rangle) \rightarrow [\exists Q\,[Q \in (S - s(M_P, w, t))]]]$

Although (34) captures the intended common-sense conception of an event in progress, it is not yet compositionally transparent. In particular, it does not make clear which ingredients are contributed by the progressive operator and which by the lexical predicate. A more perspicuous compositional picture emerges when we recognize that, as a consequence of identifying accomplishment predicates with causal models, all elements of (34) that involve model structure must be provided by the predicate itself. Consequently, the progressive operator contributes only the requirement that the reference time situation instantiate an initiated but non-completed P-eventuality. The accomplishment predicate, via its associated model M_P, supplies the relevant notions of initiation, culmination, and termination, understood in terms of the part-whole structure of sufficient sets for the truth and falsity of C_P. Our proposal thus yields an essentially partitive theory of grammatical aspect, on which the progressive relates an eventuality predicate to a reference time situation by means of *initiation* and *non-completion*:

(35) $[\![\textsc{prog}]\!] := \lambda w \lambda t \lambda P\,.\,\mathsf{INIT}([\![P]\!], w, t) \wedge \neg\mathsf{END}([\![P]\!], w, t)]$

Both INIT and END are defined as relations between an actual situation (obtaining at a world-time pair) and the causal model M_P associated with predicate P. INIT holds just in case the situation $s(M_P, w, t)$ realizes a proper part of some culmination procedure specified by M_P. END, by contrast, holds just in case the situation realizes either a sufficient set for culmination (CUL) or a sufficient set for non-culmination (TERM). The negation of END therefore excludes both culmination and termination.[37]

(36) Let P be an eventuality predicate, w a world, and t a time.

a. $\mathsf{INIT}([\![P]\!], w, t) := \exists S\,[\mathsf{Suff}_{M_P}(S, \langle C_P, 1\rangle) \wedge [\exists Q\,[Q \in (S \cap s(M_P, w, t)) \wedge$
$\exists t' \prec t\,[\forall t''\,[t'' \in [t', t) \rightarrow Q \notin s(M_P, w, t'')]]]]]]$

b. $\mathsf{END}([\![P]\!], w, t) := \mathsf{CUL}([\![P]\!], w, t) \vee \mathsf{TERM}([\![P]\!], w, t)$, where

i. $\mathsf{CUL}([\![P]\!], w, t) := \exists S\,[\mathsf{Suff}_{M_P}(S, \langle C_P, 1\rangle) \wedge (S - s(M_P, w, t)) = \emptyset\,]$

ii. $\mathsf{TERM}([\![P]\!], w, t) := \exists S\,[\mathsf{Suff}_{M_P}(S, \langle C_P, 0\rangle) \wedge (S - s(M_P, w, t)) = \emptyset]$

(35) and (36) are intended to make (34) compositionally explicit, restating the three key conditions as *initiation, non-culmination* and *non-termination.* INIT corresponds to (34)i, while ¬END abbreviates the two remaining requirements: since END is defined as CUL ∨ TERM, ¬END is equivalent to ¬CUL ∧ ¬TERM, corresponding to (34)ii and (34)iii, respectively.

Any aspectual operator must compose with an uninflected eventuality predicate, which on the present account corresponds to a causal model. As emphasized in the previous section, the progressive does not assert the existence of a *sui generis* process entity or state (cf. Mourelatos 1978; Moens and Steedman 1988; Stout 1997); rather, it asserts that, at the relevant time, some causal pathway associated with P is underway. Under (35)–(36), an eventuality counts as *in progress* just in case some, but crucially not all, of the conditions in a sufficient set for C_P are realized, while

[37]The predicates INIT, CUL, TERM, and END bear only a superficial similarity to Parsons's (1989) Cul and Hold (see 21): the latter two predicates are analyzed as primitive properties of events, whereas the present account provides explicit formal definitions for INIT, etc. relative to the causal model associated with an event type.

no sufficient set for non-culmination has been realized. The anti-termination requirement ensures, in addition, that further progress toward C_P remains compatible with the reference time facts.

On this view, then, contrasts among grammatical aspects cannot be reduced, as on more standard accounts (see Section 2.2), to differences in the temporal relation between a single actual eventuality and the reference time. Rather, aspects differ in how much of a causal pathway made available by the lexical predicate must be realized by the reference time situation, that is, in whether the relevant situation satisfies INIT, CUL, and/or TERM. Although a full partitive theory of aspect lies beyond the scope of the present paper, the analysis developed here suggests a natural treatment of non-culminating perfective uses attested in a wide range of languages (see Martin, 2019, and references therein).[38] The relations in (36) may also prove relevant to the decompositional analysis of aspectual verbs such as *begin*, *continue*, and *stop*, discussed in the introduction. We leave that extension for future work.

A fully general semantics of the progressive would, of course, extend beyond accomplishments to cover atelic activity predicates which are also crosslinguistically compatible with progressive marking. The denotation of an activity predicate cannot involve a causal structure oriented toward culmination, and the definitions above therefore do not transfer straightforwardly. At the same time, even activities may be associated with temporal boundaries, in the sense that there may be points before which the predicate is false and after which it is no longer true. Moreover, building on ideas expressed in Dowty (1979), activities can potentially be construed as pairing some consistent input of energy (DO) with a manner specification detailing how this energy is expressed. If we translate this idea into the framework of causal models, it seems to associate activity predicates with *termination* conditions: any instantiation of such a predicate necessarily comes to an end at the point where the input of energy ceases, and progressives of such predicates cannot be true after such a point. Treating activities in this way opens a route to compatibility with both INIT and TERM as defined above, and thus to a progressive interpretation along the lines of (34)/(35). We postpone a more detailed exploration of this idea to future work.

4.2 An illustration: the dual-pathway door

We illustrate our analysis with an example adapted from Baglini and Bar-Asher Siegal (2025). Suppose we have a door which can be opened either mechanically (by turning a handle) or electronically (by pressing a button). This system, schematized in Figure 2 and example (37), specifies two distinct sufficient sets for the culmination condition of the predicate *open the door*.

The model M_{door} is defined over six propositional variables: $\{H, L, E, B, C, O\}$. H, L, E, and B are background (independent) variables in this closed system: the value of H reflects whether (1) or not (0) the door *handle* is turned by an agent or effector, L reflects whether (1) or not (0) the door is *locked*, B reflects whether or not the door-open *button* is pressed by an agent or effector, and E reflects whether or not *electricity* is flowing through the door's circuit. The dependent variables in the system are C and O: here, C reflects whether or not the opening mechanism's *circuit* is closed,

[38] The central idea is that culminating and non-culminating perfectives are unified, in opposition to progressives, by a shared END requirement, while differing within and across languages in whether they specifically require END or CUL. A strong, culmination-entailing perfective, such as the English simple past or the Hindi compound perfective (Arunachalam and Kothari, 2011; Singh, 1998), would require a situation temporally included in the reference time to satisfy both INIT and CUL. A weaker, non-culminating perfective, such as the Hindi simple perfective, would require INIT and END, and would therefore be satisfied both in culminating cases and in cases where a telic eventuality is initiated but ends prior to successful culmination (see also Altshuler, 2014; Nadathur and Filip, 2021).

and O is the variable corresponding to the predicate's culmination condition, reflecting whether or not the door is *open*. The relevant causal laws are given as structural equations in (37): the circuit in the door is closed just in case the door-open button has been pressed (f_C) and the door is open if either the handle has been turned while the door is unlocked or the electronic mechanism (requiring an unlocked door, circuit closure, and the flow of electricity) has been engaged (f_O).

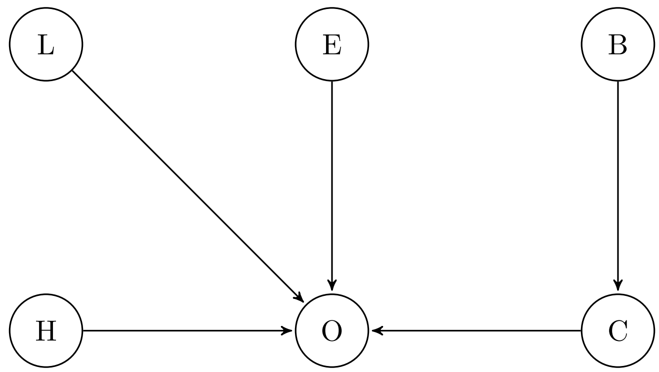

Figure 2: Graphical skeleton for the dual-route door model M_{door}

(37) a. Circuit equation (f_C): $C := B$

b. Door-open equation (f_O): $O := (H \wedge \neg L) \vee (\neg L \wedge E \wedge C)$

Table 2 provides the three-valued truth table for M_{door}, enumerating causally consistent situations according to the model; as per the structural equations in (37), the values of C and O in each line are determined by the values of their immediate causal ancestors ($Z_C = \{B\}, Z_O = \{H, L, E, C\}$).

M_{door} provides two sufficient sets (culmination procedures) for door-opening: the mechanical route corresponds to the set S_{mech} in (38a) and the electronic route to the set S_{elec} in (38b). As discussed in Section 3.3, these sufficient sets are minimal in that they comprise only necessary conditions for culmination. Any causally normal situation which contains at least the determinations in (38a) or (38b) will be one in which the door opens; this can be verified using Table 2.[39]

(38) Sufficient sets (culmination procedures) for the dual-route door to open:

a. $S_{\text{mech}} = \{\langle H, 1\rangle, \langle L, 0\rangle, \langle O, 1\rangle\}$

b. $S_{\text{elec}} = \{\langle L, 0\rangle, \langle E, 1\rangle, \langle B, 1\rangle, \langle C, 1\rangle, \langle O, 1\rangle\}$

The model also provides three sufficient sets for non-culmination, enumerated in (39). S_{lock} is realized if the door is locked, as this blocks opening in any context. The other two sets require the failure of both pathways to successful culmination: $S_{\neg\text{elec}}$ combines the failure of condition H with the absence of electricity, while $S_{\neg\text{circ}}$ combines the failure of H with the failure of button-pressing (and thus the lack of circuit closure). These sufficient sets can be verified, with a little work, by consulting the truth table in Table 2: any situation in which the door is locked ($\langle L, 1\rangle$) is one in which the door does not open, and every other line in Table 2 with $\langle O, 0\rangle$ contains the determination $\langle H, 0\rangle$ paired either with the determination $\langle E, 0\rangle$ or the pair of determinations $\{\langle B, 0\rangle, \langle C, 0\rangle\}$.

[39] B is not an immediate causal ancestor of O, but S_{elec} includes a determination for this variable precisely because this fact is *causally necessary* (by virtue of its relationship to C) for the door to open. Any change to the value of B as given in S_{elec} would require changes to the values of both C and O in order to maintain causal consistency.

H	L	E	B	C	O
u	u	u	u	u	u
u	u	u	0	0	u
u	u	u	1	1	u
u	u	0	u	u	u
u	u	0	0	0	u
u	u	0	1	1	u
u	u	1	u	u	u
u	u	1	0	0	u
u	u	1	1	1	u
u	0	u	u	u	u
u	0	u	0	0	u
u	0	u	1	1	u
u	0	0	u	u	u
u	0	0	0	0	u
u	0	0	1	1	u
u	0	1	u	u	u
u	0	1	0	0	u
u	0	1	1	1	1
u	1	u	u	u	0
u	1	u	0	0	0
u	1	u	1	1	0
u	1	0	u	u	0
u	1	0	0	0	0
u	1	0	1	1	0
u	1	1	u	u	0
u	1	1	0	0	0
u	1	1	1	1	0

H	L	E	B	C	O
0	u	u	u	u	u
0	u	u	0	0	0
0	u	u	1	1	u
0	u	0	u	u	0
0	u	0	0	0	0
0	u	0	1	1	0
0	u	1	u	u	u
0	u	1	0	0	0
0	u	1	1	1	u
0	0	u	u	u	u
0	0	u	0	0	0
0	0	u	1	1	u
0	0	0	u	u	0
0	0	0	0	0	0
0	0	0	1	1	0
0	0	1	u	u	u
0	0	1	0	0	0
0	0	1	1	1	1
0	1	u	u	u	0
0	1	u	0	0	0
0	1	u	1	1	0
0	1	0	u	u	0
0	1	0	0	0	0
0	1	0	1	1	0
0	1	1	u	u	0
0	1	1	0	0	0
0	1	1	1	1	0

H	L	E	B	C	O
1	u	u	u	u	u
1	u	u	0	0	u
1	u	u	1	1	u
1	u	0	u	u	u
1	u	0	0	0	u
1	u	0	1	1	u
1	u	1	u	u	u
1	u	1	0	0	u
1	u	1	1	1	u
1	0	u	u	u	1
1	0	u	0	0	1
1	0	u	1	1	1
1	0	0	u	u	1
1	0	0	0	0	1
1	0	0	1	1	1
1	0	1	u	u	1
1	0	1	0	0	1
1	0	1	1	1	1
1	1	u	u	u	0
1	1	u	0	0	0
1	1	u	1	1	0
1	1	0	u	u	0
1	1	0	0	0	0
1	1	0	1	1	0
1	1	1	u	u	0
1	1	1	0	0	0
1	1	1	1	1	0

Table 2: Complete truth table for M_{door}

(39) Sufficient sets for non-culmination:

a. $S_{\text{lock}} = \{\langle L, 1\rangle, \langle O, 0\rangle\}$

b. $S_{\neg\text{elec}} = \{\langle H, 0\rangle, \langle E, 0\rangle\}$

c. $S_{\neg\text{circ}} = \{\langle H, 0\rangle, \langle B, 0\rangle, \langle C, 0\rangle\}$

A truth table like Table 2 makes evident a model feature that is not immediately obvious from structural equations. A situation s corresponds to an assignment of values to variables, thus picking out a set of worlds in which all propositions corresponding to facts in $\mathsf{Ker}(s)$ are true. The truth table illustrates this space of possibilities by listing all of the relevant configurations of conditions. For the current example, Table 2 shows that the determination $\langle O, 1\rangle$ is compatible with more than one allowable configuration of M_{door}. This reflects a key causal generalization about the dual-route door—namely, that it can be opened in more than one way—which must be taken into account when the predicate *open the door* is used with this particular door in mind. Table 2 thus illustrates what it means for a single predicate to be associated with multiple type-level culmination procedures.

We emphasize that Table 2 should be interpreted epistemically. The indicated values reflect a speaker's type-level causal knowledge rather than an inventory of metaphysical facts. Consequently, a *determination* for a given variable indicates that the speaker treats the corresponding proposition as true, while a u-valuation marks the absence of settled information. Table 2 also shows that the value of a dependent variable (e.g., O) may be causally determined even when the values of one or more of its causal ancestors remain unknown. If the electrical circuit is known to be closed, for instance, the value of O may (*modulo* valuations for H and L) be settled without information about whether or not the button has been pressed; if the door is known to be locked, M_{door} supports the inference that the door remains closed regardless of the values of the other variables.

Finally, we consider how the model M_{door} contributes to the evaluation of a progressive claim. Assume that the door referenced by the definite NP in (40) is of the type specified above. Then, since M_{door} provides at least one culmination procedure for *open the door*, felicity condition (33a) is satisfied, and truth-conditional evaluation of (40) will, on our analysis, depend on a comparison between reference time facts and the sets (culmination procedures) S_{mech} and S_{elec} given in (38).

(40) Nur is opening the door.

Suppose (40) is uttered at a time when the door is known to be unlocked and Nur has just flipped the switch in the electrical cabinet, allowing electricity to flow, but has not yet pressed the button that opens the door. Such a situation realizes both $\langle L, 0\rangle$ and $\langle E, 1\rangle$, so that the reference time situation s $(= s(M_P, w, t))$ overlaps with S_{elec}, but does not share $\langle E, 1\rangle$ with the facts from just before reference time, satisfying $\mathsf{INIT}(\llbracket\text{open the door}\rrbracket, w, t)$. However, since $\langle B, 1\rangle$ has not been realized, S_{elec} is not fully realized by s; s does not yet causally determine either $\langle C, 1\rangle$ or $\langle O, 1\rangle$, so $S_{\text{elec}} - s \neq \emptyset$ and $\mathsf{CUL}(\llbracket\text{open the door}\rrbracket, w, t)$ is false, as desired. Assuming further that no sufficient set for $\langle O, 0\rangle$ has been completed—i.e., s does not fix both $\langle H, 0\rangle$ and $\langle B, 0\rangle$—$\mathsf{TERM}(\llbracket\text{open the door}\rrbracket, w, t)$ is also false, verifying (40).

The sketch above treats the variables H and B as undetermined (resulting in the absence of a causal determination for variables C and O). This is a radically epistemic view, picking out an information state corresponding to the highlighted row in the first column of Table 2. This seems somewhat artificial, however: a more realistic perspective would treat H and B as 0-determined at reference time (since our agent Nur has neither turned the door handle nor pressed the button). This scenario is picked out by the highlighted state in the second column of Table 2, which fixes the status of the door as closed ($\langle O, 0\rangle$): i.e., a state in which a sufficient set for $\mathsf{TERM}(\llbracket\text{open the door}\rrbracket, w, t)$

is in fact realized. This appears to falsify (40) relative to (34)/35 in the context described, against intuition. However, since any door-opening process must begin from a state in which the relevant door is closed, we might (more generally) think of a process for culmination as one that *changes* predetermined values of certain causally relevant variables from determinations that block culmination (support non-culmination) to determinations that are compatible with a culmination procedure, and crucially fail to converge on a procedure for non-culmination.

One way of handling this issue within the current framework is to incorporate into a causal model information about the potential fungibility of background variables. For instance, a model for the dual-route door might take into account that a lock can transition from locked to unlocked (0 to 1) as well as from unlocked to locked (1 to 0), but that a button can only transition from unpressed to pressed, so that any line of the truth table in Table 2 containing $\langle L, 1\rangle$ can become a situation where all other background variables stay the same but the value of L switches to 0 (and vice versa), whereas a line with the valuation $\langle B, 0\rangle$ can transition to a line with $\langle B, 1\rangle$, but the reverse move is not possible.[40] Building this information into the model would allow us to redefine the non-TERMination requirement (condition iii) in (34) so that a situation which verifies a sufficient set for non-culmination could still satisfy the progressive aspect as long as it permits a transition to a situation which does not verify a non-culmination procedure with at most one update to each background variable.[41] On this approach, representing the assumed context for (40) by means of the highlighted line in column 2 of Table 2 need not falsify the progressive, since the facts which currently verify $\mathsf{TERM}(s, [\![\text{open the door}]\!])$—namely, the propositions specified by $S_{\neg\text{circ}}$ in (39c)—can be updated to a non-terminating configuration by flipping either the value of H or the value of B (each of which produces a causally-triggered change in dependent variables which moves O away from a 0 valuation). A similar idea—treating certain determinations in a model as changeable while others are fixed—is explored to handle questions of agents' intentions in Lauer and Nadathur (2018); a large body of research on the lexical semantics of certain verbal predicates further supports the idea that certain changes are lexically encoded as non-reversible, while others are not (see, e.g., Bhadra, 2024, and references therein). We postpone a full exploration of this idea for future work examining the link between lexical meaning and model construction.

5 The imperfective paradox: the view from causal models

This section returns to the earlier literature on the imperfective paradox in order to clarify the broader rationale of the proposal developed here. The preceding sections introduced a causal model-based analysis of telic predicates and of the semantics of their progressive forms. Against that background, we now reconsider the central questions that motivated the discussion in the first place, and assess how they are answered once the relevant notions of process, culmination, and partial realization are reconstrued in the causal modeling perspective.

[40] This is superficially similar to the idea behind Pearl's (2000) **DO** operator, which performs 'surgery' on a causal model by setting some variable to a given value and severing links to its causal ancestors. Here, however, direct interventions are only permitted on independent variables and are further restricted to particular types of transition.

[41] This captures a key point: what matters is a contrast between 'temporary' states of non-culmination, which occur during many ordinary culmination processes, and states of non-culmination which are achieved at or before a reference interval and which force (non-reversible) termination with respect to this interval.

5.1 Revisiting the data

Our Section 2.2 discussion of intensional (Type I) approaches to the imperfective paradox focused on three types of data: progressives of impossible tasks (IEs), unlikely tasks (UEs) and telic progressives in out of reach contexts (ORCs). While normality-based approaches (e.g., Dowty 1979; Asher 1992) predict automatic falsity for all three on the grounds that the relevant reference contexts in each case fail to predict culmination, we find that IE progressives are empirically distinguished from UE and ORC progressives; see Table 1 on p.15. IE progressives are uniformly rejected, but UE and ORC progressives can be either true or false, depending on what is actually going on at reference time. Our intuition, which we used to motivate the causal approach, is that the key difference has to do with the nature of the underlying predicate: IE progressives involve tasks which are categorically unrealizable, while UE and ORC progressives do not.

By relying on an association between telic predicates and type-level causal models, the truth conditions in Section 4.1 make good on this intuition, and divide IEs from UEs and ORCs as desired. On our analysis, IE progressives like (16a)-(16b) are rejected due to infelicity, rather than strict falsity. Use of a telic predicate P requires the existence of corresponding causal model, establishing at least one cognitively-plausible procedure for realizing culmination condition C_P: a predicate that describes a task known to be impossible is, *ipso facto*, a predicate for which the type of world knowledge that comes into play in the construction of an event type causal model litigates against the existence of such a pathway. In the absence of a model, there is no object against which reference time facts can be evaluated in the terms set out in Section 4.1, and IE progressives are therefore not truth-conditionally evaluable.

Unlike IEs, UEs and ORCs both involve predicates that pick out events that are possible (if non-trivial) to fully realize, and which thus have event type models as far as the typical language user is concerned. Progressives of these predicates can simply be true or false of a given situation, depending on the validity of all three conditions in (34)/(35); there is no meaningful difference between UEs and ORCs on the causal approach. In the example of the undertrained runner, for instance, (19a) (see p.11) will be true as long as Benny has started the race (e.g., by taking a step into the race path after the starting gun has fired; INIT), has not reached the end of the race path ($\neg$CUL), and is able at reference time to continue towards completion ($\neg$TERM). Crucially, the truth of the progressive does not take into account the likelihood that Benny will complete the race: once he has initiated an appropriate process, all that matters at a given moment is that it remains possible for him to make further progress. This possibility collapses as soon as Benny's stamina runs out, but not before: in the terms in (34)/(35), the total depletion of energy which accompanies his inevitable collapse completes a sufficient set of conditions for non-culmination, (irreversibly) validating TERM (36b-ii) and thus falsifying the $\neg$END requirement of (35).

5.2 Partitive treatments for the imperfective paradox

The contrast between Type I and Type P approaches to the imperfective paradox turns on the role of intensionality. Type I approaches preserve the culmination assumption by deriving 'paradoxical' effects from modality introduced by the progressive; Type P approaches, by contrast, provide a partitive and extensional treatment of aspect, by allowing non-culminated eventualities to instantiate telic predicates. We have seen that the Type I strategy encounters empirical difficulties as a result of tying the truth of the progressive to locally projected culmination. At the same time, the Type P alternative leaves open a key question: what determines whether a given reference time situation qualifies as a partial realization of a telic predicate, rather than as a merely related activity? We argued in Section 2 that, once the puzzle is recast in partitive terms, one must still specify the (abstract) whole relative to which actual events can be construed as *proper parts*.

Our proposal adopts a broadly partitive view, since it involves evaluation of the part-whole relation between actual events and the causal pathways that constitute the denotation of a telic predicate. It is not, however, an extensional account. Although events that fall short of culmination can validly instantiate a telic type, they do so only by virtue of correspondence to a culminating whole defined by a type-level causal model. To the extent that such a model encodes generalizations about how events of a given type are typically realized, it is an inherently intensional object. One effect of our proposal, then, is to relocate the modality associated with non-culmination: rather than coming from aspectual modification, it is built into the denotations of telic predicates themselves (see also Koenig and Muansuwan 2000; Fiorin and Delfitto 2017; Nadathur and Filip 2021).

Since causal models encode information about how a culmination condition is ordinarily realized, our proposal retains a connection to the notions of normality and stereotypical development that motivate Type I approaches. The difference lies in where this information enters the analysis and in the theoretical work it must do. On a Type I view, normality is determined relative to the evaluation context, and progressive truth depends on whether the reference time situation supports a local projection of culmination. On our view, by contrast, normality is anchored in the predicate itself, through the type-level causal structure associated with its lexically specified culmination condition. The relevant question is therefore not whether the reference time situation itself predicts C_P, but whether it instantiates a recognizable stage in a causally normal pathway to C_P. Put differently, the issue is whether the situation corresponds a sort of cross-section—i.e., a situation that can arise at some stage in the development of an unexceptional culminated P-eventuality.[42] Shifting from token-level expectations of culmination to type-level causal structure is what allows the present account to handle UE and ORC progressives: a reference-time situation may instantiate such a stage even when, in the context of evaluation, culmination is not expected to occur.

The point is illustrated by the dual-route door in Section 4.2. The model M_{door} associated with the predicate *open the door* is not a model of any particular (token) event, but rather a type-level representation of how a door of this sort is opened, and thus a way of encoding a range of possible door-opening events that may be partially or fully instantiated in a world. The context discussed for (40)—i.e., a situation in which an agent has ensured that electricity is flowing but has neither turned the door handle nor pressed the door-open button—verifies the progressive because it matches a developmental cross-section of at least one unexceptional culminated door-opening eventuality: for instance, one in which the electric route to door-opening (S_{elec}; 38b) is fully deployed. Example (40) is true in this context because part of a type-level causal pathway is realized in real time.

In taking a partitive but still intensional view, our proposal has something in common with Type I analyses that take the internal structure of complex events into account. Landman (1992), for example, introduces a primitive notion of *event stages* (wherein "an [actual] event is a stage of another [actual] event if the second can be regarded as a more developed version of the first"; p.23) in order to treat the progressive as a relation between an actual event e and a predicate P: $\text{PROG}(P, e, w)$ is true just in case e is a stage of a (culminated) P-eventuality in the **continuation**

[42] This is reminiscent of teleological modality, which is concerned with what is necessary or possible in worlds where some relevant goal or *télos* is realized (see, e.g., von Fintel and Iatridou 2005). A causal model for predicate P can in this sense be understood as characterizing a set of idealized, causally normal worlds in which the *télos* C_P is realized. See Nadathur and Filip 2021 for further development of this connection.

branch of e with respect to w. The crux of the account is the notion of a continuation branch, which tracks event e across a set of 'reasonable' alternatives to w: we follow e to the point at which it ceases in w, and if this is due to some event-external interruption, we shift attention to the continuation of e in the nearest reasonable alternative where the interruption does not occur. The process is repeated until the tracked event either comes to its natural culmination continuation moves beyond the 'reasonability' horizon.

One advantage of continuation branches over the Type I proposals discussed in Section 2 is that interruptions need not be treated as abnormal in any objective sense: for Landman, $\text{PROG}(P, e, w)$ holds just in case there is a reasonable chance that the processes defining an actual event e will continue to develop towards C_P at each point at which they might be interrupted. UE progressives can then be verified because the continuation branch 'removes' potential interruptions sequentially (Landman, p.30)—all that is required is that each potential point of disruption (however likely it may be) might plausibly be avoided, allowing the event to continue for a little bit longer. The appeal to reasonability also rules out IE progressives like (16a), as desired: even if Mary gets into the ocean at Galway (event e), and even if there is one reasonable alternative in which she continues swimming westwards, drawing on a last reserve of energy, past the point at which she actually succumbs to exhaustion, doing so at every potential failure point will presumably take us past the reasonability horizon well before the crossing is complete, so that the continuation branch of e cannot verify the progressive with respect to the target predicate.[43]

(16a) #/??Mary is swimming across the Atlantic Ocean.

Like our analysis, Landman's approach considers actual events as potential developmental stages of full P-eventualities, and prioritizes the possibility of continuation over a modal expectation of culmination. This allows him to make the right predictions for UE progressives. Insofar as the continuation branch must culminate within the set of reasonable options (projected from world w and event e), Landman's PROG is ultimately an existential intensional operator. We saw Section 2.2 that this approach cannot explain acceptable ORC progressives: indeed, it is difficult to see how Landman's explanation for the failure of IE claims—that culmination simply crosses the reasonability horizon—would not also apply to the case of the undertrained runner in (19). Reliance on the existence of an underlying causal model (to provide the truth conditions for a valid event stage) allows us to make the necessary distinction between IEs and ORCs, where reliance on what is reasonable on the basis of local facts does not.

[43] Landman (1992) discusses the need to balance the "problem of interruptions" with what he calls the "problem of non-interruptions", referring to a shift in judgments observed in case an ostensibly impossible task is actually completed. In such a scenario—e.g, where we learn from a credible witness that Mary (having started to swim west from Galway some time ago) has just climbed out of the ocean in Newfoundland—the past tense correlate of (16a) becomes true, even though (16a) would not have been accepted at the same reference time. Standard Type I approaches struggle to account for this type of temporal asymmetry, but since a continuation branch first follows an actual event e to its maximal development, the fact that Mary has actually succeeded in swimming across the Atlantic cannot be ignored: the culminating continuation branch is fully contained in the evaluation world, and what actually happens is enough to verify the (past) progressive.

We have previewed how the causal perspective handles this data. The underlying intuition is similar to Landman's: whatever actually happens constitutes a course of events that is, *ipso facto*, causally normal. We assume that the actual world is treated as causally normal (i.e., that causal reasoning employs as a default that "anything that has happened must be possible"), so that knowledge of Mary's success forces the conclusion that some sufficient set of causal conditions for culmination was actually realized, and thus that some culmination procedure exists: it is true that *Mary was swimming across the Atlantic* because the actual world demonstrates not only that there is a way to do so, but also that Mary was (at reference time), engaged in the right sort of process.

Bonomi (1997) offers another Type I analysis that incorporates the internal structure of complex events. His underlying intuition is that an actual event can be seen as a part (or stage) of a range of distinct, more developed eventualities, but that which—if any—of these larger events is realized depends on the context in which the reference event occurs. On Bonomi's account, the progressive expresses an intensional relation between an actual event e and a predicate P by projecting a set of "courses of events" (akin to Thomason's 1984 *histories*) from e, taken together with some contextual set C_H of facts. These courses of events are *stereotypical* in the sense that they permit e to follow its natural development: $\textsc{prog}(e, P)$ is true relative to C_H just in case all of the "ideal" developments of e which conform to C_H contain an event f of type P such that e is a stage of f.

UE progressives can be true for Bonomi because the alternatives projected by PROG are relativized to the reference event e. He argues that the "ideal" courses of events exclude "all external factors that might block the completion of the event" (p.190), taking into consideration only what has actually happened thus far as well as concomitant facts that have to do with direct participants in the named event. This is similar to our treatment of UE progressives in prioritizing information about what is actually going on at reference time, but is differentiated by a selective reliance on context—in this case, on positive, supposedly event-external information which supports the idea that successful culmination is *possible* (if not especially likely).

This context-dependence makes it difficult to extend Bonomi's account to a successful treatment of ORC cases: it is not clear what contextual information should count as event-internal and event-external but relevant as a concomitant fact as opposed to a potential interruptor.[44] Returning to case (19) of Benny the runner, excluding his (under)preparedness from the relevant set of concomitant facts seems to allow the truth of the target progressive (19a), as desired; however, it is not at all clear that this information can be innocently excluded from C_H given the context in (19). Indeed, reference to concomitant facts of this sort suggest that Bonomi ultimately relies on the idea that some course of events in the stereotypical set is a possible future with respect to the local context in which e is embedded. To make matters worse, the processes in which Benny is involved after the race begins very reasonably include a process of energy depletion, so that even stereotypical continuations of the reference event e should reasonably prevent it from being seen as part of a larger, ultramarathon-completing eventuality. By fully divorcing notions of event structure from the local context—and by evaluating the progressive as a relation between reference time facts and an event-type model—the causal approach simplifies the calculation here: we do not need to consider the relative trajectories of actual ongoing processes, but simply whether not a set of instantaneous facts (Benny's intentions, his placement along the race course, that he has at least some energy remaining) correspond to some developmental stage in an idealized race-completion.

The intensional but event-structural proposals developed by Landman and Bonomi may, in principle, be extended to capture the contrast between IE and UE/ORC cases, provided that notions such as 'reasonable', 'appropriate', or 'sufficient' partial realization are taken to bear substantial

[44] Bonomi (1997) motivates introducing event mereology by observing that neither a requirement of temporal nor of causal precedence (between the actual event e and P's culmination) alone is enough to justify the progressive: for instance, while all events which culminate in a painted house must be both temporally and causally preceded by some agent's acquisition of the necessary painting materials, an event in which the agent purchases equipment will not in general verify the progressive with respect to house-painting. Introducing an atomic notion of event parthood pushes this issue into a question of world knowledge: while it might not be straightforward to specify general conditions on the activities A which constitute an initiated part of a predicate P, common sense judgments supply information on a case-by-case basis. This is the type of information that is presumably encoded in a causal model, or—for Bonomi—which govern the part-whole relations between events in a stereotypical set of histories.

explanatory weight. From this perspective, one might hope to recover the relevant distinctions by appealing to a suitably constrained notion of what counts as a legitimate stage of a telic eventuality. Our view, however, is that such notions remain insufficiently explicit as explanatory primitives. The contribution of the present paper is therefore not merely to preserve the descriptive contrast, but to provide a framework in which these informal notions are replaced by objects that admit independent formal characterization, namely, causal structures defined over event types. Whereas the accounts discussed briefly here rely on a comparatively underspecified mereological notion of event *stages*, our proposal models developmental (process) components of a telic eventuality type in terms of a structured set of necessary conditions that are jointly sufficient for the realization of the culmination condition C_P. On this approach, then, the question of when a reference time situation qualifies as a genuine partial realization of a telic predicate receives a precise and theoretically grounded answer.

5.3 Reference time structure as the basis for partial realization

The account with the closest resemblance to ours is that of Varasdi (2014). As we did in Section 2.1, Varasdi uses ORC examples like (19) to argue that the truth of a progressive claim with telic predicate P does not require that culmination is a possibility from the reference time perspective. Instead, he treats the progressive as a property of states which can be true of a reference situation s just in case s verifies conditions which are necessary for a (culminated) event of type P. The analysis is motivated by the asymmetric entailment relationship between a (culminating) telic perfective and its corresponding progressive (see Section 2), which (cf. Szabó 2004) shows that the progressive of telic P is verified by a strict subset of the conditions that validate its perfective counterpart.

Varasdi proposes that a context in which $\text{PROG}(P)$ is used implicitly makes available a set R of distinct relevant outcomes which includes (the culmination condition of) P. The progressive claim is true in such a context just in case there is a particular facet F (subset of facts) of the reference situation s which is *indicative* of P in R—i.e., such that the characteristics of s which are picked out by F are only compatible with P. This is operationalized via three conditions in a possible-worlds framework: first, that some world w which realizes the facet F of s is such that some mereological development of s (i.e., some event of which s is a part) in w is an event of type P; secondly, that no such world realizes an alternative outcome from the set R as a mereological development of s; and finally, that worlds which do not verify F as a facet of s are worlds in which P is not realized, making the conditions in F necessary (albeit not necessarily sufficient) for P.

Although Varasdi does not discuss IE progressives directly, the account should treat these as false (not infelicitous), provided that the set of worlds from which the perspective verifying F as a facet of s is drawn are restricted to worlds which are broadly causally normal: under this assumption, while the perspective picked out by s and F may well be non-empty, it cannot contain a full realization of the IE predicate, and thus the first condition for a telic progressive fails. By virtue of being achievable in causally plausible albeit atypical circumstances, UE progressives can be satisfied by a state of affairs just in case there is some perspective which verifies a set of necessary conditions for the underlying predicate P which uniquely select from the set of contextually-relevant outcomes. IE progressives work much the same way: as above, we need to abstract away from local facts—such as Benny's ill-preparedness—which preclude the right sort of culmination, but this exclusion can occur as part of the selection of a given facet/perspective on the reference situation, and need not be justified by considerations of what is internal/external to an ongoing reference time event, as on Bonomi's (1997)'s account.

The conceptual resemblance between Varasdi's proposal and ours is immediate. Both analyses prioritize the information included in a reference time state of affairs over the contextually-expected continuations of this state, with the result that UE and ORC progressives can be true. In both cases, the idea that the evaluation context must realize *sufficiently much* of an event of type P is reduced to a type of subset relation: as Varasdi notes, necessary conditions are (can be) sufficient. However, while both accounts require that the reference situation contains some set of facts which are necessary for a full realization of predicate P, Varasdi's reliance on alethic necessity (i.e., strict entailment) implicitly requires that there is only one way—one set of sufficient conditions—which can fully realize a given predicate P. This may be appropriate for predicates with relatively homogeneous process components (*run a marathon, climb the stairs*), but seems less appropriate for predicates whose culmination conditions may be realized in a variety of distinct ways (such as the dual-route door example in Section 4.2); defining necessary conditions with respect to their role in potentially distinct sufficient sets in an event type causal model, as we do in Section 3.3, allows for this possibility. It may of course be possible to amend Varasdi's analysis to handle this concern, but—to the extent that his account must independently incorporate a mereological event structure according to which a given state can be seen as part of a larger eventuality—our view is that the causal modeling approach streamlines implementation of the shared underlying ideas.

Summarizing Sections 5.1-5.3, the causal model approach to telic progressives allows us to make the notion of partial realization—with respect to an abstract but conceptually real whole—explicit, and thereby to capture in a principled way the contrast between IE and the other key data types discussed above. The crucial distinction is not simply whether culmination occurs, or even whether it is likely to occur in the circumstances at hand, but whether the predicate is associated with a causally coherent culminatory structure at all. Moreover, because a causal model for a telic predicate encodes not only satisfaction conditions but also explicit non-satisfaction conditions, it is possible to distinguish genuine non-culmination from cases in which culmination remains in principle available but is improbable. The framework defines model structure over event types rather than event tokens, separating contingent failures in a concrete individual cases from the existence of a process whose culminatory structure remains well defined. This is what makes it possible to treat ORC cases as valid uses of the progressive, even when an event under consideration cannot (from the local perspective) culminate. More broadly, the causal modeling approach provides an explicit theoretical vocabulary, drawn from independent notions in causal reasoning, for articulating distinctions that have often remained intuitive or loosely characterized in the previous literature. What emerges is not merely a new treatment of telic progressives, but a more general perspective on how event structure should be represented. If the analysis developed here is on the right track, the imperfective paradox can be seen as a diagnostic for a broader fact: the interpretation of telic predicates depends on access to structured, type-level causal knowledge.

6 Further theoretical issues

This section moves beyond comparisons with existing literature to illustrate additional theoretical consequences of our analysis. On the account developed here, telic predicates are associated with type-level causal models, and progressive predication is understood in terms of partial realization of a causally structured path toward culmination. Once this perspective is adopted, two questions arise. The first concerns the epistemic position from which a relevant causal model is selected and employed. The second concerns the status of particular conditions—including agentive intention—which play a role in structuring and guiding an event towards culmination. By showing how these issues fit within the present framework, we aim to further clarify the interaction between world knowledge, model selection, and the interpretation of telic predicates.

6.1 Epistemic perspective and model selection

Insofar as the content of a causal model depends on world knowledge, a language user's epistemic state plays a non-trivial role in the evaluability and the evaluation of telic progressives. The use of a telic predicate presupposes the availability of an appropriate causal model in at least the following sense: it signals speaker commitment to the cognitive availability of such a model, as well as the expectation that the model is shared with their interlocutors, at some relevant level of granularity.[45]

IE progressives are infelicitous because no suitable shared model is available.[46] Thus, an audience confronted with an example such as the hole-digging claim in (16b) is unlikely to accommodate the speaker's apparent presupposition that there is some way for children (or for anyone) to dig a hole to China, precisely because, for a well-informed adult, the achievability of the task is patently absurd. Once the speaker's epistemic perspective is made salient, however, the judgment changes:

(41) **The revised hole-digging scenario.** Meena's five-year-old daughter Maya falsely believes that the earth is made entirely of sand and soil and that its diameter is much less than it actually is. Maya is digging a hole at the beach, with the stated intention of tunneling all the way through the planet.

a. *Maya:* I am digging a hole to China!

b. *Meena:* ??Maya is digging a hole to China.

In (41), Maya's epistemic state—which includes false beliefs—supports a relatively simple strategy for digging a hole to China, and thus the existence of a plausible type-level causal model for the task, relative to her own perspective on the actual world. Although Maya is, of course, objectively mistaken, and although it would be inappropriate for an adult observer to sincerely describe her activity by means of (41b), it does not seem right to classify (41a) as a lie. Relative to Maya's own information state, her utterance is both appropriate and true.

The causal modeling approach allows us to capture this contrast directly. A well-informed adult observer will not treat Maya's self-description as authoritative with respect to objective reality, but neither will that observer typically regard it as pragmatically deviant. The reason is that the interpretation of the progressive depends, in part, on the body of world knowledge relative to which the relevant causal model is constructed, and this body of knowledge may vary across speakers. The framework therefore predicts a limited form of perspective dependence: the felicity and truth of a telic progressive may shift with the epistemic state that anchors the relevant causal model, even where speakers agree about the observable facts.

6.2 Globally necessary conditions in the structure of telic events

The causal-model approach developed so far treats telic predicates as encoding structured sets of pathways leading to a culmination condition. However, not all components of these pathways are

[45] We assume that, when a hearer encounters a novel telic predicate, the default strategy is to treat the speaker's information state as *prima facie* authoritative. Thus, in the absence of background knowledge that rules out the existence of an appropriate model—as in IE cases—an unfamiliar predicate is likely to be accepted as felicitous.

[46] McHugh (2023) derives a related restriction from the fact that a model determines, for each possible world, whether the relevant outcome is nomically possible or impossible. In our approach, as noted above, the restriction follows more directly from the nomological character of the model itself.

on equal footing: some are contingent features of particular procedures, while others appear to be required across all realizations of the event type. Conditions that are **globally necessary** in this sense (GNCs) appear to play a special role in the evaluation of telic predicates.

Consider once again the ultramarathon scenario in (19). Benny's attempt to run the race unfolds along a recognizable pathway toward culmination, but certain conditions—in particular, his intention to participate in the specified race—plays a crucial role in structuring this development. Assuming that running an official race is an agentive (volitional) pursuit, such an intention is plausibly required throughout any process that leads to successful completion: Benny's bid to run an ultramarathon may terminate unsuccessfully before his energy is fully depleted if he simply decides he no longer wants to participate. Extrapolating from this, an agent's intention to realize the culmination condition C_P of an agentive predicate acts as GNC for culmination: a condition that must hold across all causal pathways leading to C_P, and which sustains their development towards this point (cf. Varasdi 2014 on *sustaining conditions*). Formally, GNCs can be modeled as propositions whose singular falsity guarantees non-culmination, as spelled out in (42):

(42) **Globally necessary conditions.** Fact Q is *globally necessary* for C_P with respect to model M_P just in case $\mathsf{Suff}_{M_P}(\{\neg Q, \langle C_P, 0\rangle\}, \langle C_P, 0\rangle)$.[47]

GNCs are worth discussing because their determinations (0-1 valuation) appear to constitute minimal requirements for the truth-conditional assessment of a telic progressive, even if certain other conditions which are plausibly part of a causal pathway for culmination have already been realized.

The GNC effect can be illustrated with the dual-door context from Section 4.2: here, the state of the lock functions as a GNC, since the door cannot be opened when the lock is engaged. Suppose that the status of the lock is unknown ($\langle L, u\rangle$), and that Nur has either turned the door handle or pressed the door-open button. In neither situation is it felicitous to assert (43a); the strongest claim that can be truthfully made is (43b). This suggests that the lock condition is not simply one contributing factor among many, but is one which must be resolved before the relevant event can be verified or disqualified as a door-opening process: this is the case even if Nur has performed some other action that would normally initiate door-opening.

(43) a. ??Nur is opening the door.
b. Nur is trying to open the door.

The same effect can be illustrated for agentive predicates by varying the running context:[48]

[47] We use $\neg Q$ to represent the ordered variable-value pair that shares the variable coordinate of a fact Q but flips its value between 0 and 1. Since Q corresponds to a proposition, this notation reflects standard logical negation.

[48] Example (44) is an instance of a *multiple-choice scenario* (Bonomi, 1997) As Bonomi points out, Landman's (1992) continuation branch analysis makes the wrong predictions for contexts of this sort. The causal modeling approach can appeal to GNCs to explain why a disjunctive progressive can be true in a context in which neither disjunct is acceptable on its own—context (44) satisfies an intention GNC for (44c) but not for (44a) or (44b)—but, for Landman, the truth of (44c) requires that one of (44a) or (44b) is true: the culmination branch that satisfies the disjunctive progressive must verify a exactly one of a completed 10 mile run or a completed half marathon.

Varasdi's (2014)'s account (discussed in Section 5.3 may also face problems here, since it relies on the idea that necessary conditions are only enough to verify the progressive of a predicate P when these conditions uniquely indicate P from a set of contextually-relevant outcomes. Pragmatically speaking, it seems reasonable that the individual options $\sqrt{}$*Benny run 10 miles* and $\sqrt{}$*Benny run a half marathon* are included in the relevant set out outcomes whenever $\sqrt{}$*Benny run 10 miles or a half marathon* is; however, if all three alternatives are included, it should not be possible to select a facet of any reference time situation which admits events verifying the disjunctive outcome without also admitting events verifying one or the other of the individual disjuncts.

(44) **The variable distance scenario.** Benny began running in a race which had different distance options: 10 miles, half marathon, and full marathon, with separate finish lines along the main course. Uncertain about his precise level of preparedness, Benny did not intend to complete the full marathon distance, but planned to decide at the 10 mile turnoff whether to stop there or continue to run a half marathon. He unexpectedly collapsed between the 8th and 9th mile marker, before making a decision. Later, he says:

a. ??I was running 10 miles when I collapsed.

b. ??I was running a half marathon when I collapsed.

c. ✓ I was running 10 miles or a half marathon when I collapsed.

Since Benny has the disjunctive intention to run 10 miles or a half marathon and has taken steps along the race path, (44c) is true: up to the moment of collapse, he was on a pathway for culmination (for one of these tasks) which was neither completed nor terminated. However, since he lacks both the specific intention to run exactly 10 miles and the intention to run a half marathon, neither (44a) nor (44b) seems to be an appropriate description of his activities. We can capture the judgments in (43a) and (44a)-(44b) by constraining our analysis of telic progressives as in (45).[49]

(45) **Determination of GNCs.** PROG(P) is defined at $\langle w, t\rangle$ iff no GNCs for C_P are undetermined in context:

$$[\![\text{PROG}(P)]\!]^{w,t} \in \{0,1\} \leftrightarrow \forall Q\, [\mathsf{Suff}_{M_P}(\{\neg Q, \langle C_P, 0\rangle\}, \langle C_P, 0\rangle) \rightarrow \{Q, \neg Q\} \cap s(M_P, w, t) \neq \emptyset]$$

In both agentive and non-agentive cases, GNCs seem to involve properties which 'guide' a P-eventuality towards culmination. For a non-agentive accomplishment predicate with a homogeneous process component—e.g., *roll off the table*—a reasonable candidate for a GNC (analogous to agentive intent) might be a conserved quantity like momentum. Knowledge of such quantities—unlike knowledge of intentions, which must be inferred from agents' behaviour or accepted based on self-report—is a matter of fine-grained physical observation, to a degree which may not be possible for the casual observer. This leads to an interesting predicted contrast between agentive and non-agentive progressives. On our proposal, Benny's underdetermined intention in (44) is sacrosanct: there is simply no truth of the matter with respect to (44a) or (44b) until and unless he makes up his mind. In a parallel non-agentive case, such as (46), the progressive of one of the disjuncts is ostensibly true, but its determination is out of reach of the standard (unaided) human observer.

(46) A fair coin is tossed, and is still rising.

a. ?The coin is coming up heads.

Bonomi avoids the multiple-choice problem by incorporating a modal and context-sensitive mereological structure for event types (p.188). This explains the judgments in (44) if we take the reference event to comprise Benny's race-related activities, including information about his point of origin, his progress thus far, his speed, energy, heart rate, and so on. This event could be seen as part of any number race experiences, but its development is constrained by local ("concomitant") facts (see Section 5.2) such as the amount of pre-race fuel Benny has ingested, how hot it is, and any intentions he has with respect to the length of his race. Including his intention to run either 10 miles or a half marathon in the constraining set restricts the set of stereotypical developments of Benny's race thus far to include attempts which end either at the 10 mile or half marathon mark, validating (44c) but falsifying (44a) and (44b) (since the contextual facts do not successfully pick out a single distance).

[49] This is a conservative form of the GNC constraint. We suspect that (45) is not peculiar to telic progressives, but instead represents a restriction on the INIT condition defined in Section 4; however, we have yet to fully investigate the relationship between GNCs and non-progressive uses of telic predicates.

b. ?The coin is coming up tails.

(47) a. *In (44):* Benny is running a 10 mile race or a half marathon, ??but I don't know which.

b. *In (46):* The coin is coming up heads or tails, but I don't know which.

This prediction—and thus the causal approach to telic progressives—gains some preliminary support from the contrasting judgments in (47), but a detailed investigation is left for future research.

Finally, an interesting and positive consequence of (45) is a straightforward account of a surprising temporal asymmetry in the truth and felicity of telic progressives, first noted by Bonomi (1997): past tense progressives can sometimes be verified with respect to a reference time at which the corresponding present tense claim could not have been truthfully uttered. For instance, if we modify the scenario in (44) so that Benny reaches the 10 mile mark and decides to stop there, the past tense claim in (48a) becomes true for any reference time between the time at which he started the race and the time at which he stopped running; (48b) would not have been true if uttered at any of those times. Unlike most Type I approaches to telic progressives, the GNC presupposition in (45) predicts this asymmetry, since a GNC (here, Benny's distance-relevant intention) was not determined until the 10 mile mark. GNCs, in other words, represent necessary conditions for the conceptual initialization of a process for C_P, but their determination need not correspond to the temporal initiation of such a process.

(48) a. Benny was running 10 miles.

b. Benny is running 10 miles.

The asymmetry can also emerge in another way: we find that the empirical judgments of (48a)-(48b) are the same if Benny reaches the 10 mile marker and collapses without making a concrete decision about his goals. This suggests, at least at first glance, that agents' intentions are relevant to judgments of telic progressives only from epistemic perspectives at which the outcome of the event in progress remains undetermined—i.e., in the present tense, but not necessarily in the past.

This observation suggests an alternative way of understanding the role of intentions in these cases—one that does not rely on their characterization as GNCs. On this view, the contrast reflects not a difference in the truth conditions themselves, but an epistemic constraint on the identification of the relevant causal model. Accomplishments such as running a particular distance can be realized both as deliberate, *agentive* acts and as non-agentive processes, and these may correspond to distinct models. When Benny completes a 10 mile run and then stops unintentionally, the model for the agentive form of the predicate becomes irrelevant: his prior actions can be identified directly as belonging to a realization of a non-intentional culmination pathway for a 10 mile run, and the relevant *télos* (and thus the relevant model) is fixed retrospectively by the actual course of events. While the process is still unfolding, however, truth-conditional judgments of the progressive depend on the ability to identify a specific model for comparison. In scenarios that admit multiple potential outcomes (see also Bonomi 1997 on the *problem of underdetermination*), evidence about an agent's intentions serves as a cue for selecting among competing models. In the absence of such evidence—whether because it is underdetermined or simply inaccessible to the speaker—the issue is not that the truth conditions fail, but that they cannot yet be evaluated, since no determinate predicate model is available for comparison.

This also explains why one of the past tense correlates of (46a)-(46b) becomes true, and the other false, as soon as the coin has landed. The difficulty with the present tense claims is not due to the failure of the underlying conditions, but arises from the epistemic inaccessibility of fine-grained facts that determine the outcome while the event is ongoing. These facts can only be inferred retrospectively, once the event has culminated; at that point, it becomes possible to say, of the earlier time when the coin was still in the air, that it *was coming up heads* or *was coming up tails*.

This discussion also bears on the epistemic aspects of our proposal. We regard causal models as language-independent cognitive devices for representing how events normally unfold: i.e., as a means of encoding world knowledge about the conditions under which certain types of event can develop. One epistemic element is incorporated through the use of a tripartite truth value system, allowing us to reason about situations where the actual value of one or more conditions is unknown; Section 6.1 showed that there are contexts in which telic predicates are evaluated relative to individual speakers' epistemic states. In this section, we see that assessment of a telic progressive (or, extrapolating, any use of an inflected telic predicate) may also depend on what kind of epistemic evidence is available at a given point in time. In real-time (present tense) evaluation, independent evidence concerning agents' intentions and other GNCs appears crucial for determining whether an actual situation can qualify as part of a relevant culmination process. In some retrospective perspectives, however—specifically, when culmination has taken place—independent evidence is no longer required: the existence of a culminated event itself licenses the classification of its earlier parts as stages of a process leading to culmination (cf. Szabó 2004; see also fn.43).

7 Conclusions and outlook

We began by identifying a lexical puzzle for durative telic predicates: how the clear and intuitive connection between such predicates and their lexically-specified culmination conditions can be reconciled with their robust use in non-culminating contexts. Using the imperfective paradox as a case study, we argued that the core theoretical challenge is to provide a satisfactory notion of *partial realization*, specifying formal conditions under which a non-culminated eventuality realizes *sufficiently much* of a telic event type to validate the predicate. In the progressive domain, these are precisely the conditions under which an actual event counts as making progress toward culmination.

Although telic progressives have been our primary focus, we began by discussing bipartite theories of lexical decomposition in order to highlight certain general issues that underlie our analysis. First, decompositional theories foreground the question of what belongs to the lexical semantic content of the predicate itself. One of the main conclusions of the present study is that the denotation of a telic predicate must include idiosyncratic content pertaining not only to culmination (as is frequently assumed), but also with respect to the structured processes that leads to culmination. On our view, process and culmination content are both represented (and linked) by a type-level causal model.

Second, decompositional approaches bring into focus the role of causation itself, including the question whether a predicate like CAUSE should be treated as a semantic primitive and whether causal structure is adequately understood as a simple binary relation between two events. Although our proposal links telic predicates to inherent causal structure, it departs in a fundamental way from classical causative analyses of the telic aspectual class, which have been strongly contested in the literature (see, e.g. Levin and Rappaport Hovav, 1999; Levin, 2000). On standard CAUSE-decompositional approaches, some causing event involving the sentential subject stands in a uniform binary causal relation to the actual realization of the culmination condition, with the result that the underlying predication of actual causation can be validated only when both cause and effect are actualized in the evaluation context. Our proposal retains the insight that telic predicates invoke causal structure, but relocates that structure to the level of event types. Actualized process components of telic predicates are not instantiated *actual causes*, but causally organized assemblies of conditions that belong to sufficient sets for culmination. The details of such processes, as well as the relations among their components, are supplied by world knowledge organized formally within a model. Use of a telic progressive therefore conveys commitment to the existence of a cognitively plausible causal model for the underlying predicate, but does not itself predicate an actual causal relation to culmination: the result is that a telic predicate can be associated with causal relationships without its felicitous use requiring the realization of both (some) cause and its effect. Similarly, realization of a valid process eventuality does not reduce to a localized modal prediction of culmination, but instead requires a match between reference time facts and some developmental stage in a type-level pathway for culmination.

Finally, standard bipartite decompositional theories bring into focus what we refer to in this paper as the *culmination assumption*: the idea that telic predicates denote exclusively culminated eventualities. This assumption follows directly from the bipartite approach and lies at the heart of the problem of telic non-culmination; while our discussion focuses on the resultant imperfective (progressive) paradox, it should be clear that the culmination assumption raises problems that extend beyond progressive contexts. As suggested in Section 4, we believe that adopting the causal modeling perspective for telic predicates makes it possible to provide parallel accounts of non-culmination phenomena that arise in the composition of telic predicates with other aspectual operators, such as non-culminating perfectives and aspectual verbs.

Looking further ahead, we suspect that telic predicates are not unique in admitting a fruitful analysis in terms of causal models. If the present proposal is on the right track, then the success of causal modeling in the progressive domain should be understood not merely as a local solution to one aspectual puzzle, but as evidence for a more general rethinking of event structure. The broader program suggested here is one which understands predicates of eventualities as labels for structured bodies of causal information: collections of conditions that enable, sustain, constitute, or otherwise lead to particular states or outcomes. On such a view, familiar aspectual class properties—including durativity, telicity, and related distinctions (Vendler, 1957; Kenny, 1963; Dowty, 1979, a.o.)—need not be stipulated independently, but may instead be derived from properties of the relevant causal models. The proposal advanced in this paper is therefore intended not only as an account of telic progressives, but also as a first step toward a broader theory in which causal models form part of the core of event conceptualization and event-structural representation: that is, toward an aspect calculus based on the structural properties of formal causal models, as suggested in the introduction.

Declarations.

Conflict of interest. The authors have no relevant financial or non-financial interests to disclose.

References

Abusch, Dorit. 1985. *On Verbs and Time*. Ph. D.thesis, University of Massachusetts, Amherst.

Alonso-Ovalle, Luis and Hsieh, Henrison. 2021. Causes and expectations: on the interpretation of the Tagalog ability/involuntary action form. *Journal of Semantics* 38, 441–472.

Altshuler, Daniel. 2014. A typology of partitive aspectual operators. *Natural Language and Linguistic Theory* 32, 735–775.

Arunachalam, Sudha and Kothari, Anubha. 2011. An experimental study of Hindi and English perfective interpretation. *Journal of South Asian Linguistics* 4, 463–508.

Asher, Nicholas. 1992. A default, truth conditional semantics for the progresive. *Linguistics and Philosophy* 15, 463–508.

Bach, Emmon. 1986. The algebra of events. *Linguistics and Philosophy* 9, 5–16.

Baglini, Rebekah and Bar-Asher Siegal, Elitzur. 2025. Modelling Linguistic Causation. *Linguistics and Philosophy* 48, 647–691.

Baglini, Rebekah and Francez, Itamar. 2016. The implications of managing. *Journal of Semantics* 33(3), 541–560.

Bar-Asher Siegal, Elitzur. 2025. Causation and causal relations: insights from natural language. *Annual Review of Linguistics* 12, 125–145.

Bar-Asher Siegal, Elitzur, Bassel, Noa and Hagmayer, York. 2021. Causal selection: the linguistic take. In *Experiments in Linguistic Meaning*, volume 1, pages 27–38.

Bar-Asher Siegal, Elitzur and Boneh, Nora. 2020. Causation: from metaphysics to semantics and back. In Elitzur Bar-Asher Siegal and Nora Boneh (eds.), *Perspectives on Causation: selected papers from the Jerusalem 2017 workshop*, pages 3–51, Cham, Switzerland: Springer.

Baratella, Riccardo. 2023. Processes and their modal profile. *Synthese* 201, 74.1–24.

Beavers, John and Koontz-Garboden, Andrew. 2012. Manner and result in roots of verbal meaning. *Linguistic Inquiry* 43, 331–369.

Bhadra, Diti. 2024. Verb roots encode outcomes: argument structure and lexical semantics of reversal and restitution. *Linguistics and Philosophy* 47, 557–610.

Bjorndahl, Adam and Snider, Todd. 2015. Informative counterfactuals. In Sarah D'Antonio, Mary Moroney and Carol Rose Little (eds.), *Semantics and Linguistic Theory*, volume 25, pages 1–17.

Bonomi, Andrea. 1997. The progressive and the structure of events. *Journal of Semantics* 14, 173–205.

Ciardelli, Ivano, Zhang, Linmin and Champollion, Lucas. 2018. Two switches in the theory of counterfactuals: a study of truth conditionality and minimal change. *Linguistics and Philosophy* 41, 577–621.

Copley, Bridget. 2021. Reconciling modal and causal representations of Salish out of control forms. In *West Coast Conference on Formal Linguistics*, volume 39, pages 495–503.

Copley, Bridget and Wolff, Philip. 2014. Theories of causation should inform linguistic theories and vice versa. In Bridget Copley and Fabienne Martin (eds.), *Causation in Grammatical Structures*, pages 11–57, Oxford: Oxford University Press.

Davidson, Donald. 1967. The logical form of action sentences. In Nicholas Rescher (ed.), *The Logic of Decision and Action*, pages 81–120, Pittsburgh: University of Pittsburgh Press.

Dowty, David. 1977. Towards a semantic analysis of verb aspect and the English imperfective progressive. *Linguistics and Philosophy* 1, 45–77.

Dowty, David. 1979. *Word Meaning and Montague Grammar*. Dordrecht: Reidel.

Dowty, David. 1991. Thematic Proto-Roles and Argument Selection. *Language* 67, 547–619.

Fiorin, Gaetano and Delfitto, Denis. 2017. A Perspective-Based Account of the Imperfective Paradox. In Joanna Blochowiak, Cristina Grisot, Stephanie Durrleman and Christopher Laenzlinger (eds.), *Formal Models in the Study of Language: Applications in Interdisciplinary Contexts*, pages

97–121, Cham: Springer International.

Gehrke, Berit and McNally, Louise. 2009. Frequency adjectives and assertions about event types. In Ed Cormany, Satoshi Ito and David Lutz (eds.), *Semantics and Linguistic Theory*, volume 19, pages 180–197.

Grimshaw, Jane. 2005. Semantic structure and semantic content in lexical representation. In Jane Grimshaw (ed.), *Words and Structure*, pages 75–89, Stanford, CA: CSLI Publications.

Hale, Ken and Keyser, Samuel Jay. 2002. *Prolegomenon to a Theory of Argument Structure*. Cambridge, MA: MIT Press.

Halpern, Joseph Y. 2015. A modification of the Halpern-Pearl Definition of Causality. In Qiang Yang and Michael Wooldridge (eds.), *International Joint Conference on Artificial Intelligence (IJCAI)*, volume 24, pages 3022–3033, Palo Alto, CA: AAAI Press.

Halpern, Joseph Y. and Pearl, Judea. 2005. Causes and explanations: A structural-model approach – Part I: Causes. *British Journal of Philosophy of Science* 56, 843–887.

Hausman, Daniel M. 1998. *Causal Asymmetries*. New York: Cambridge University Press.

Hausman, Daniel Murray. 2005. Causal Relata: Tokens, Types, or Variables? *Erkenntnis* 63, 33–54.

Henderson, Robert. 2014. If not for counterfactuals: negating counterfactuality in natural language, ms, University of Arizona.

Hitchcock, Christopher. 2020. Communicating causal structure. In Elitzur Bar-Asher Siegal and Nora Boneh (eds.), *Perspectives on causation: selected papers from the Jerusalem 2017 workshop*, Jerusalem Studies in Philosophy and History of Science, pages 53–71, Springer Cham.

Hobbs, Jerry. 2005. Toward a useful concept of causality for lexical semantics. *Journal of Semantics* 22, 181–209.

Hoover, Kevin. 2001. *Causality in Macroeconomics*. Cambridge: Cambridge University Press.

Hume, David. 1748. *An Enquiry Concerning Human Understanding*. Oxford: Oxford University Press.

Kaufmann, Stefan. 2013. Causal premise semantics. *Cognitive Science* 37, 1136–1170.

Kenny, Anthony. 1963. *Action, Emotion, and Will: Studies in Philosophical Psychology*. London: Routledge and Kegan Paul.

Klein, Wolfgang. 1994. *Time in Language*. New York: Routledge.

Koenig, Jean-Pierre and Muansuwan, Nuttanart. 2000. How to end without ever finishing: Thai semi-perfectivity. *Journal of Semantics* 17, 147–182.

Kratzer, Angelika. 1981. The Notional Category of Modality. In Hans-Jürgen Eikmeyer and Hannes Rieser (eds.), *Words, Worlds, and Contexts. New Approaches in Word Semantics*, pages 38–74, Berlin: de Gruyter.

Kratzer, Angelika. 1996. Severing the external argument from the verb. In Johan Rooryck and Laurie Zaring (eds.), *Phrase Structure and the Lexicon*, pages 109–137, Dordrecht: Kluwer.

Kratzer, Angelika. 1998. More structural analogies between pronouns and tense. In Devon Strolovich and Aaron Lawson (eds.), *Semantics and Linguistic Theory*, volume 8, pages 92–110, Ithaca, NY: CLC Publications, Cornell University.

Kratzer, Angelika. 2004. Telicity and the meaning of objective case. In Jacqueline Guéron and Jacqueline Lecarme (eds.), *The Syntax of Time*, pages 389–423, Cambridge, MA: MIT Press.

Krifka, Manfred. 1989. Nominal reference and quantification in event semantics. In Renate Bartsch, Johan van Benthem and Peter van Emde Boas (eds.), *Semantics and Contextual Expression*, pages 75–115, Dordrecht: Foris Publications.

Landman, Fred. 1992. The progressive. *Natural Language Semantics* 1, 1–32.

Lascarides, Alex. 1991. The progressive and the imperfective paradox. *Synthese* 87, 401–447.

Lauer, Sven and Nadathur, Prerna. 2018. Sufficiency causatives, ms, Universität Konstanz and Stanford University.

Levin, Beth. 2000. Aspect, lexical semantic properties, and argument expression. In *Berkeley Linguistics Society*, volume 26, pages 413–429.

Levin, Beth and Rappaport Hovav, Malka. 1991. Wiping the slate clean: a lexial semantic exploration. *Cognition* 41, 123–151.

Levin, Beth and Rappaport Hovav, Malka. 1999. Two Structures for Compositionally Derived Events. In Tanya Matthews and Devon Strolovitch (eds.), *Semantics and Linguistic Theory*, volume 9, pages 199–223.

Levin, Beth and Rappaport Hovav, Malka. 2005. *Argument realization*. Cambridge: Cambridge University Press.

Levin, Beth and Rappaport Hovav, Malka. 2013. Lexicalized meaning and manner/result complementarity. In Boban Arsenijević, Berit Gehrke and Rafael Marín (eds.), *Subatomic Semantics of Event Predicates*, pages 49–70, Dordrecht: Springer.

Levin, Beth and Rappaport Hovav, Malka. 2014. Manner and result: a view from *clean*. In Rob Pensalfini, Myfany Turpin and Damian Guillemi (eds.), *Language Description Informed by Theory*, pages 337–357, Amsterdam: John Benjamins.

Lewis, David. 1973a. Causation. *Journal of Philosophy* 70, 556–567.

Lewis, David. 1973b. *Counterfactuals*. Oxford: Blackwell.

Lewis, David. 1979. Counterfactual dependence and time's arrow. *Noûs* 13, 455–476.

Link, Godehard. 1983. The logical analysis of plurals and mass terms. In Rainer Bäuerle, Christoph Schwarze and Arnim von Stechow (eds.), *Meaning, Use, and Interpretation of Language*, pages 302–323, Berlin: Walter de Gruyter.

Mackie, John. 1965. Causes and conditions. *American Philosophical Quarterly* 2, 245–265.

Martin, Fabienne. 2018. Time in probabilistic causation: direct vs. indirect uses of lexical causative verbs. In *Sinn und Bedeutung*, volume 22, pages 107–124.

Martin, Fabienne. 2019. Non-culminating accomplishments. *Language and Linguistics Compass* 13, 1–20.

McHugh, Dean. 2023. *Causation and Modality*. Ph. D.thesis, University of Amsterdam.

Moens, Mark and Steedman, Mark. 1988. Temporal ontology and temporal reference. *Computational Linguistics* 14, 15–28.

Mourelatos, Alexander. 1978. Events, processes, and states. *Linguistics and Philosophy* 2, 415–434.

Nadathur, Prerna. 2023a. *Actuality inferences: causality, aspect, and modality*, volume 15 of *Oxford Studies in Semantics and Pragmatics*. Oxford: Oxford University Press.

Nadathur, Prerna. 2023b. Causal semantics for implicative verbs. *Journal of Semantics* 40, 311–358.

Nadathur, Prerna and Filip, Hana. 2021. Telicity, teleological modality, and (non-)culmination. In *West Coast Conference on Formal Linguistics*, volume 39, pages 617–625.

Nadathur, Prerna and Lauer, Sven. 2020. Causal necessity, causal sufficiency, and the implications of causative verbs. *Glossa,* 5, 49.1–37.

Parsons, Terence. 1989. The progressive in English: events, states, and processes. *Linguistics and Philosophy* 12, 213–241.

Parsons, Terence. 1990. *Events in the Semantics of English: A Study in Subatomic Semantics*, volume 19 of *Current Studies in Linguistics*. Cambridge, MA: MIT Press.

Partee, Barbara. 1973. Some structural analogies between tenses and pronouns. *The Journal of Philosophy* 70, 601–609.

Pearl, Judea. 2000. *Causality*. Cambridge: Cambridge University Press.

Ramchand, Gillian. 2008. *Verb meaning and the lexicon: a first-phase syntax*. Cambridge: Cambridge University Press.

Reichenbach, Hans. 1947. *Elements of Symbolic Logic*. New York: The Free Press.

Schulz, Katrin. 2011. If you'd wiggled A, then B would've changed. *Synthese* 179, 239–251.

Singh, Mona. 1998. On the semantics of the perfective aspect. *Natural Language Semantics* 6, 171–199.

Sosa, Ernest and Tooley, Michael. 1993. *Causation*. Oxford: Oxford University Press.

Spirtes, Peter, Glymour, Clark and Scheines, Richard. 1993. *Causation, Prediction, and Search*. New York: Springer Verlag.

Steward, Helen. 2013. Processes, Continuants, and Individuals. *Mind* 122, 781–812.

Steyvers, Mark, Tenenbaum, Joshua, Wagenmakers, Eric-Jan and Blum, Ben. 2003. Inferring causal networks from observations and interventions. *Cognitive Science* 27, 453–489.

Stout, Rowland. 1997. Processes. *Philosophy* 72, 19–27.

Szabó, Zoltan. 2004. On the progressive and the perfective. *Noûs* 38, 29–59.

Szabó, Zoltan. 2008. Things in progress. *Philosophical Perspectives* 22, 499–525.

Thomason, Richmond. 1984. Combinations of tense and modality. In Dov Gabbay and Franz Guenther (eds.), *Handbook of Philosophical Logic: Extension of Classical Logic*, pages 135–165, Dordrecht: Reidel.

Tooley, Michael. 1987. *Causation: A Realist Approach*. Oxford: Oxford University Press.

Varasdi, Karoly. 2014. Making progressives: necessary conditions are sufficient. *Journal of Semantics* 31, 179–207.

Vendler, Zeno. 1957. Verbs and times. *The Philosophical Review* 66, 143–160.

Verkuyl, Henk. 1972. *On the Compositional Nature of the Aspects*. Dordrecht: Reidel.

Vlach, Frank. 1981. The semantics of the progressive. In Philip Tedeschi and Annie Zaenen (eds.), *Syntax and semantics 14: Tense and aspect*, pages 271–294, New York: Academic Press.

von Fintel, Kai and Iatridou, Sabine. 2005. What to do if you want to go to Harlem, ms, Massachusetts Institute of Technology.

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

Zucchi, Sandro. 1999. Incomplete events, intensionality, and imperfective aspect. *Natural Language Semantics* 7, 179–215.