Events are the source of causal readings in the simplest English conditionals

Bridget Copley, SFL (CNRS/Paris 8)

Abstract The distinction between stative predicates such as *be the princess* and eventive predicates such as *become the princess* is highly relevant to the distribution of causal readings of conditionals. To investigate this connection, here I examine the morphologically simplest conditionals in English: namely, those with no aspectual, modal, future, tense, or counterfactual morphology on their prejacents. I argue that the relationship between eventivity and causal readings of these conditionals can be understood using causal models as in Pearl (2000), along with attention to two points: first, not all seemingly eventive verb phrases are truly eventive; and second, eventive predicates are predicates of relations. Taking these points into consideration, it appears that causal relations within the meaning of eventive verb phrases are the source of causal readings of the simplest English conditionals.

1 A causal puzzle about statives and eventives in conditionals

Suppose that there is a country called Clavarel where the queen is selected from amongst the eligible candidates by a series of coin tosses. In the particular situation we are interested in, it so happens that if the coin comes up heads, Yolanda becomes the queen; and if the coin comes up tails, she does not. Suppose also that if someone becomes the queen, any daughters she may have consequently become princesses; moreover, there is no other way to become a princess of Clavarel. Further suppose that Yolanda has exactly one daughter, named Xanthippe.

In such a context, we can see that (1a) is true, while (1b) is false. What's wrong with (1b)? It certainly seems like it conveys a causal relation which is at odds with what we know—it's backwards in fact. (1b) seems to say that if Xanthippe becomes the princess, then *by consequence* Yolanda becomes the queen. (1a), on the other hand, has no such flavor.

(1) a. If Xanthippe is the princess, Yolanda is the queen.
b. If Xanthippe becomes the princess, Yolanda becomes the queen.

The puzzle here is, why? Taking seriously Bennett's comment that "The search for a deep understanding of conditionals would go better if the taxonomy were got right first" (Bennett 1988), I will make two methodological choices in teasing out (merely) where to start with the taxonomy.

1.1 Causation rather than time

The first choice has to do with time. In addition to feeling causally backwards, (1b) feels temporally backwards; it seems to say that Xanthippe becomes the princess of Clavarel *first*, and only then does Yolanda become the queen; contrary, again, to what we know about Claverelian dynastic politics. Now, the distinction between stative predicates (e.g., *be the* *princess*) and eventive predicates (e.g., *become the princess*) can be cast in terms of times; see, for example, Hallman (2009). We might therefore consider recruiting the concept of time to try and solve this puzzle. This is certainly a plausible move.[1]

However, I am going to avoid making reference to time in explaining the contrast in (1). There are two reasons for this. One reason to start from causation is that temporality can be in part derived from causation from the fact that causes don't happen after their effects, so that taking causation as primitive gets us some temporality for free as it were. On the other hand, taking time as primitive gets us nothing in the way of causal meaning. A second reason is that eventivity and stativity are phenomena internal to the verb phrase; approaches to the verb phrase across various frameworks have concluded that much, if not all, of eventive verbal meaning is organized in terms of causal relations (e.g. Dowty 1979, Pustejovsky 1995, Ramchand 2008, Croft 2012, Copley & Harley 2015). For these reasons, I will defer discussion of times and instead start from causation, relying on a simple (structural equation) version of causal models (Pearl 2000).[2]

A causal model is a directed acyclic graph used to represent causal influences (Pearl 2000, Pearl & Mackenzie 2018). For example, our causal knowledge about the situation in Clavarel can be represented as the following, where C is whether the coin turns up heads, P is whether Xanthippe is the princess, and Q is whether Yolanda is the queen.

(2) $C \longrightarrow Q \longrightarrow P$

There can be different interpretations of causal models like the one above (notably, the probabilities of the nodes are often used). In the interpretation we will use, we will use simply the nodes' truth values.

(3)
C := 1 if heads, 0 if tails
Q := 1 if Yolanda is the queen, 0 if Yolanda is not the queen
P := 1 if Xanthippe is the princess, 0 if Xanthippe is not the princess

Supposing that Yolanda becomes queen upon the coin's landing one way or another, regardless of whether the Clavarelian nobles have seen which side the coin has landed on, the table in (4) shows the combinations of truth values that are possible given the causal information in (3).

(4)

	C	Q	P
a.	1	1	1
b.	0	0	0
c.	1	0	0
d.	1	1	0

1 The fact that antecedents like *If Xanthippe becomes princess* are grammatically present but conceptually future-oriented was noticed by Dudman (1983), if not earlier; see also Kaufmann (2004) and Kaufmann et al. (2006). The significance of this fact is still controversial to philosophers, see Edgington (1995, 2003), Bennett (2003: p. 15).

2 Linguists are beginning to make use of causal models; see Kaufmann (2002, 2005, 2013), Baglini & Francez (2016), Nadathur (2016), Lauer & Nadathur (2018), Baglini & Bar-Asher Siegal (2019), Bar-Asher Siegal & Boneh (2019)

Each line of the table is a possible *situation*; I assume in this very limited scenario that 1 and 0 are the only possible values and that each node has a value in each situation. It should be uncontroversial that lines (a) and (b) are possible in our Clavarel scenario, but those in (c) and (d) may or may not be possible, depending on further details. The idea is that, even though there is strictly speaking no time between the coin coming up heads and, e.g., Yolanda becoming queen, it is still possible to be in a state where one has had the thought that the coin has come up heads but not yet concluded that Yolanda is now queen, as in line (c). However, don't get hung up on the issue of whether lines (c) and (d) belong in the table or not. The puzzle, and our solution to the puzzle, are the same regardless of whether they are there or not. In section 4, we will look at a more concise way to generate the lines of the table according to the arrows in the causal model graph, which will hopefully shed light on this issue.

1.2 Compositionality and the stative-eventive distinction

So, the first methodological choice was to use causal models. The second methodological choice is to take a compositional perspective. Semantic composition allows us to theorize about unknown meanings of parts of a sentence if we know the meaning of the whole. In this, it is a kind of algebra. And, as in algebra, it is helpful to start with the simplest expressions in order to solve for unknowns. This is why we are using the morphologically simplest prejacents in (1) to understand how *if* works; we will not even be looking at conditionals with *will* in the consequent.

Very quickly we see that the relevant distinction between (1a) and (1b) is one of stativity (as in *being*) and eventivity (as in *becoming*). The distinction between eventivity and stativity is perhaps the major division in the lexical aspect (or "Aktionsart") of verb phrases. We might think of eventive predicates as processes going on in time, with phases succeeding one another; states on the other hand, at least as far as we will talk about them here, can be thought of as "a single, undifferentiated period" (Smith 1991: 19). Using stative predicates, as in (1a), there is no causal reading, and using eventives, as in (1b), there is a causal reading. We would like to know why this is so.

1.3 *Do* will not do

Pearl and Mackenzie's (2018) discussion around causal models suggests a possible way to link event semantics to causal models. They make a distinction between "seeing" and "doing". For example, one should distinguish *If I see that Xanthippe is the princess, I see that Yolanda is the queen* and *If I make Xanthippe the princess, I make Yolanda the queen.* Their point is that seeing or inferring correlations between values of variables is not at all the same as changing or "wiggling" the value of those variables. Only the latter gives us truly causal information, in the way that turning a light switch on and off tells us which appliance it is connected to.

This idea enjoys pride of place in "interventionist" theories of causation (e.g. Woodward 2006); for Pearl, it also marks the difference beween talking about correlations and talking about causation at all (Pearl & Mackenzie 2018: Chapter 1). They note that in the field of statistics, there has been a real reluctance to use causal language for lack of a formalized way to talk about it. And indeed it is illegitimate to talk about causation using only tools that express correlation. However, as they point out, once we have a formal way to talk about causation, it is no longer illegitimate to do so.

Pearl's 2000 formal tool for talking about causation is the "*do*-operator". The *do*-operator erases all causal influences on the node in question and changes the value of the node to the specified value, as shown in (5):

(5) $do(X = x)$: erase all incoming arrows into X and change the value of X to x.

Looking at the difference between seeing and doing in this system, one could wonder whether this distinction might correspond to the distinction in English between statives and eventives respectively. After all, seeing is stative and doing is eventive. It would be very nice if this system developed for other uses could turn out to explain the stative/eventive distinction and thereby explain the contrast in (1).

The causal model in (2), together with treating conditionals as strict, works well enough for (1a). That is, given the model and the possible values, it is true that all cases where Xanthippe is the princess of Claverel, Yolanda is the queen. There is one situation where Xanthippe is the princess, namely the first line of the table in (4), and on that line, Yolanda is queen.

But what we would like to know is whether the *do*-operator, "off the shelf," as it were, might help us understand the meaning of eventive predicates *become the princess* or *become the queen* as in (1b). For suppose it did. Then if the conditional is again treated as strict implication, the meaning of (1b) would be expected[3] to be as follows:

(6) All situations where $do(P(x) = 1)$ are situations where $do(Q(y) = 1)$

That is, all cases where we erase the arrows into the $P(x)$ node and set the value of the $P(x)$ node to 1 are cases where we erase the arrows into the $Q(y)$ node and set it to 1.

But even then, we do not explain our puzzle. We do get (1b) to be false, but for the wrong reason. Making the reasonable assumption that we have the freedom, as erasers of arrows and setters of values, to apply the *do*-operator to whatever node we want, it is indeed false, as desired, that any case where we $do(P(x) = 1)$ is a case where we $do(Q(y) = 1)$. But although we get the right result, this reason for getting that result is exactly wrong. The problem with (1b) is not that there is no arrow between $Q(y)$ and $P(x)$; it is that there *is* an arrow going into $P(x)$ in the causal model, but somehow the conditional seems to require an arrow of causal influence to go the other way. In short, the *do*-operator cannot help us understand here how eventive predicates relate to causal readings in conditionals, because it erases one of the very things that we need to use, namely that incoming arrow.

1.4 How we will solve the puzzle

The *do*-operator is not the only possible way to model "doing," i.e. intervention. We can retain something like the intuitive idea that *seeing* and *doing* are to be distinguished, and

3 To underline what we are doing here: (6) is merely *our* expected meaning if we are treating each eventive verb as involving a *do*-operator, not Pearl's expectation. For one, he does not distinguish between eventive and stative verbs.

that this distinction matters for statives and eventives. I will rely on another tradition to make sense of this idea, namely dynamic semantics (Heim 1982, Kamp 1981, Groenendijk & Stokhof 1990). More closely hewing to our characterization of the puzzle in (1), statives (as in *being*) and eventives (as in *becoming*), will be treated as distinct modes of updating the world. Asserting a proposition in dynamic semantic frameworks changes the context it is uttered in; propositions are "context change potentials". Statives, we will see, behave like propositions. I will argue here that to account for the behavior of eventive predicates, we need them to be, in effect, "world change potentials". Crucially, learning or accepting a proposition does not require it to be false (or for that matter true) before the learning or accepting, whereas an event's happening does put a condition on the initial state; for Xanthippe to become a princess, she has to not already be a princess.

The analogy between propositions as context change potentials and events as world change potentials suggests that we might want to model eventives as predicates of ordered pairs of atoms, not predicates of atoms. That is, although in the neo-Davidsonian tradition we typically think of causal relations between subevents, here the events will actually be themselves represented by an ordered pair between two situations. The types[4] we will be interested in are in (7):[5]

(7) a. stative predicate: $\langle s, t\rangle$
b. eventive predicate: $\langle s{\times}s, t\rangle$

Accordingly, there will be two ways for the speaker to update the current situation: *learning*, which is an update with a predicate of situations, and *happening*, which is an update with a predicate of relations between situations.

(8) a. learning = update with a predicate of situations
b. happening = update with a predicate of a relation between situations

These two kinds of updates, associated with the verb phrases themselves, will be situated with a dynamic treatment of *if* φ, ψ.

Working from these components, I will argue that the event in the consequent is, literally, the causal relation that yields the causal reading in (1b). To argue this, first we will set out to understand what eventives are, both in terms of the tests we use to distinguish eventives and statives, and in the sense of how they appear in conditonals. It turns out that not all apparent eventives in conditional antecedents are really, truly, eventives. The ones that are not get a stative reading, so they must be excluded if we want to say something about the behavior of eventives. These can be thought of as eventives with a stative component on top of them, so they may not be truly semantically "bare" eventives even though they are

4 As is usual in formal semantics, I will assume that the denotations of phrases are functions from typed objects to other typed objects, where these typed objects are perhaps themselves such functions. Type is a formal property, which allows us to compose meanings with each other according to a small collection of compositional rules, including Functional Application Heim & Kratzer (1998). In short, if two denotations are not of the right type to compose, they will be unable to compose. The basic types used here are s, the type of situations, and t, the type of truth values.

5 The type of the eventive predicate could also be written as $\langle\langle s, s\rangle, t\rangle$, as in Copley & Harley (2015, 2021). The (equivalent) cross notation used here underlines the fact that the relation between situations, as far as type theory and semantic composition are concerned, is not treated like a compound type.

morphologically bare. This point relates to our methodological decision to consider only the simplest conditionals, namely those with no aspectual, modal, future, tense, or counterfactual morphology on the prejacents. Recall that this choice is an important one, as it allows us to see how event semantics interacts with conditional semantics when there is nothing semantically active (in the relevant sense) insulating the verb phrase (the locus of event semantics) from *if* (the locus of conditional semantics).

Subsequently, we will see evidence from assertability facts to the effect that "true" eventives (i.e., "truly" bare eventives in the sense we have seen) have a type difference from statives and larger phrases with aspect, modals, futures, and/or tense. I will use these to argue for the type difference in (7) above, where events take us from one situation to another in much, but not entirely, the same way that utterances do in dynamic semantics.

My proposal to solve the puzzle in (1) is given in section 4, with three components: (a) the causal model is used to create a causal sequence of situations, (b) eventives update the situation using a relation between situations instead of a situation, and (c) *if* φ, ψ is provides sequential updates. With this framework, we see that the combination of the eventive predicate in the consequent and the causal model itself ultimately provides the causal reading of the conditional and thereby explains the puzzle.

A last note before we dive in. This line of inquiry is nowhere near the point of proposing a theory of all conditionals. It turns out to be quite enough for now to argue that event semantics and causal models can be linked to account for the puzzle.

2 True eventives are a different type than statives

2.1 Tests for eventivity vs. stativity

To understand what statives and eventives are doing in the simplest English conditionals, we have to have tests to distinguish statives from eventives. Such tests are notoriously language-specific and can be difficult to interpret, but two of the most robust tests are shown in (9) and (10). The first test involves *must*. In (9a), which has a stative predicate, the epistemic reading of the *must* sentence is available; with an eventive, as in (9b), the epistemic reading is not available.

(9)	a.	Xanthippe must be sick.	epistemic reading possible ⇒ stative
	b.	Xanthippe must get sick.	no epistemic reading possible ⇒ eventive

The second test is variously referred to as a biscuit conditional, an Austinian conditional, or a relevance conditional (Austin 1961, DeRose & Grandy 1999, Krifka 2014: and references therein). The idea behind this test is that there can be no causal connection between the prejacents; rather, there is a kind of discourse relation (which one it is has been argued rather closely but is not important to its use here as a test for stativity). So, (10a) comes out something like "if you care to know, then it is relevant to you/I assert that there is beer". To get this reading on a simple conditional it is required to have a stative in the consequent, as in (10a). Using an eventive in the consequent, as in (10b), does not permit this reading.[6]

6 That is, it does not permit this reading as long as we are not dealing with a generic or habitual reading; this is exactly the point to be discussed just below.

(10) a. If you care to know, there's beer. relevance reading ⇒ stative consequent
b. If you care to know, we get beer. no relevance reading ⇒ eventive consequent

While there is agreement in formal semantics that a broad eventivity/stativity distinction is crucially visible to a number of grammatical phenomena, mostly related to aspect and modality, there are various ways to analyze the distinction. Common to many of them is Davidson (1967)'s argument that the meanings of sentences such as *Jones buttered a piece of toast* make reference to an abstract event argument. In further "neo-Davidsonian" developments of this idea, beginning with Higginbotham (1983, 1985, 1986) and Parsons (1990), we can think of the verb phrase "butter the toast" as predicating buttering the toast of an event variable e. This idea can be extended to states; how far to extend it is a live issue of contention (see, for instance Kratzer 1995, Maienborn 2005, 2007). For now the assumption we will make is that the verb phrases, whether intuitively eventive or stative, are to be treated as predicates. The question at issue will be what they are predicates of. We will see that while statives make sense as predicates of situations, it makes sense for eventives to be predicates of relations between situations.

2.2 Some apparent "eventives" in English are actually stative

It turns out that very often in English, predicates that we expect to be eventive actually behave like statives instead. These stative readings of apparently eventive predicates fall into several categories: generic/habitual readings, futurate readings (where there is a future-oriented, "planned" or "settled" reading but no future morphology, (Lakoff 1971, Vetter 1973, Dowty 1979, Kaufmann 2005, Copley 2008, 2018)), and what I will call "storytelling readings", where the reference is to what happens in a story, play, etc. (cf. the "director's reading" in Ritter & Rosen (1997)). The fact that generic/habitual readings are stative is well-known, while the behavior of futurate and storytelling readings has not been much talked about (but see Copley (2018)).

So, for instance, an "eventive" predicate such as *drink beer* permits a generic/habitual and epistemic reading for (11a) in a context where one is looking around the throne room and one sees lots of empty beer bottles. Likewise, (11b) gets a relevance reading.

(11) Generic/habitual readings behave like statives
a. Clavarelian nobles must drink beer.
b. If you care to know, Clavarelian nobles drink beer.

Futurate readings behave exactly the same way. A context that favors the futurate reading in (12a) is one where, for instance, we see on Xanthippe's calendar "Z: Green Linnet" for tomorrow, and we know that her best friend is Zelinda and the Green Linnet serves only beer. In this context an epistemic reading is possible, and we might also felicitously use the relevance conditional in (12b).

(12) Futurate readings behave like statives
a. Xanthippe must drink beer with Zelinda tomorrow.
b. If you care to know, Xanthippe drinks beer with Zelinda tomorrow.

Finally, storytelling readings behave exactly the same way, permitting epistemic *must* as in (13a) and (when in the consequent) relevance conditionals as in (13b).

(13) Storytelling readings behave like statives

a. (I haven't read to the end of the book, but from what I've read,) Abelard must drink beer at the end of the book.

b. If you care to know, Abelard drinks beer at the end of the book.

We can ask whether stativity is really the property that makes futurates and storytelling examples behave like statives—indeed, we will return to this question. But it is incontrovertible that in (12) and (13) respectively, futurate and storytelling readings do behave like statives.

So, several readings of "eventive" predicates behave like statives do with respect to the stativity tests introduced above. Readings that are not generic/habitual readings, not plannable (so do not license futurate readings), and not set in a story (so do not have storytelling readings), do not behave this way.

(14) "None-of-the-above" readings do not behave like statives

a. Yolanda must get sick tomorrow. — no epistemic reading

b. #If you care to know, Yolanda gets sick tomorrow.

(15) a. It must rain tomorrow. — no epistemic reading

b. #If you care to know, it rains tomorrow.

What follows from this is that habitual/generic readings, futurate readings, and storytelling readings are just statives; or more precisely, that the highest predication in their structure is stative. This could happen in one of (at least) two ways. Either a "derived stative" method, where they have an unpronounced stative operator that takes the eventive predicate as an argument; or a *sui generis* method, where there never is truly an eventive predicate in the semantics and the root itself, or something about the predication itself, is allowed to pick out a stative meaning, and there is still morphological "bareness", i.e. no extra operator. Both possibilities have been suggested for generic/habitual readings: either there is a Gen operator that takes the eventive predicate as an argument, or there is simply a stative predication of something like a kind argument (see Krifka et al. (1995), Carlson (2009). For futurates, many, including Copley (2018), have proposed that futurate meaning is supplied with an unpronounced operator; Kaufmann (2005) is unique as far as I know in leaving the door open for the other possibility. There is a very interesting question here about what exactly morphology tells us about the complexity of denotations; it is immaterial, however, to the present paper.

Returning to the context where a fair coin toss determines whether or not Yolanda becomes the queen, we see that *become the queen* behaves like a (true) eventive, not a stative. This predicate is episodic, cannot be planned so is not futurate,[7] and it is not—to the people living it—in a story.

7 Futurates require temporal adverbials; if (and only if) we drop the assumption that it is a fair election and add a temporal adverbial, we do indeed get a futurate reading. In that case *Yolanda must become the queen tomorrow* has a possible epistemic reading and *If you care to know, Yolanda becomes queen tomorrow* is a good relevance conditional.

(16) "None-of-the-above" readings do not behave like statives (to a speaker within the story, so not a storytelling reading)

a. Yolanda must become the queen.
b. #If you care to know, Yolanda becomes the queen.

Just to remember why the above facts are important: Recall that I am arguing that the causal reading of (1b) depends on its consequent being eventive. It is important to the argument therefore that we only consider truly eventive readings.

2.3 True eventives are not assertable

Now that we can distinguish true eventives from merely apparent eventives, we can start to argue that statives and true eventives have a type difference. The first step is to see that true eventives in English are not assertable.

(17) True eventive

a. If it rains tomorrow...
b. #It rains tomorrow.

(18) True eventive

a. If Yolanda gets sick tomorrow...
b. #Yolanda gets sick tomorrow.

This observation has been noted by Dudman (1983), Edgington (2003), a.o. Of course, in contrast to (17) and (18), many of the simplest antecedents are assertable on their own, as shown in (19) through (22) below. These correspond, again, to statives, generic/habitual, futurate, and storytelling readings, as in (19), (20), and (21), and (22) respectively:

(19) a. If Xanthippe is there...
b. Xanthippe is there.

(20) a. If Zelinda drinks beer...
b. Zelinda drinks beer.

(21) a. If Xanthippe drinks beer with Zelinda tomorrow...
b. Xanthippe drinks beer with Zelinda tomorrow.

(22) a. If Abelard drinks beer at the end of the book...
b. Abelard drinks beer at the end of the book.

Finally, while in antecedents, true eventives are possible, it seems that in the consequents of the simplest conditionals (recall, these are conditionals with no aspectual, modal, future, tense, or counterfactual morphology on the prejacents), as long as we resist the temptation to add *will* or similar, true eventives are not possible. Instead, we always get a generic/habitual, a futurate or a storytelling reading. Adding *will* is what makes the episodic reading possible. For example, there is no reading of (23) below which means that if Zelinda should happen to drink beer, Xanthippe will also happen to drink beer. Instead, we get generic/habitual readings either low, in each prejacent, or high, scoping over the whole conditional:

(23) Generic/habitual reading
(If Zelinda drinks beer,) Xanthippe drinks beer.

a. 'If Zelinda habitually drinks beer, Xanthippe habitually drinks beer.
b. 'Generally, if Zelinda drinks beer, Xanthippe drinks beer."
c. #'If Zelinda happens to drinks beer, Xanthippe will happen to drink beer.'

The same is true for futurate readings; a low and a high scope for the plan are both possible, but without *will*, there's no merely episodic, non-futurate reading.

(24) Futurate reading
(If Zelinda drinks beer tomorrow), Xanthippe drinks beer tomorrow.
a. 'If there's a plan for Zelinda to drink beer tomorrow, there's a plan for Xanthippe to drink beer tomorrow.'
b. 'There's a plan such that if Zelinda drinks beer tomorrow, Xanthippe drinks beer tomorrow.'
c. #'If Zelinda happens to drink beer tomorrow, Xanthippe will happen to drink beer tomorrow.'

(24a) and (24b) are possible, but what we emphatically do not get is an episodic reading as in (24c) where, Zelinda's happening to drink beer tomorrow has something to do with Xanthippe happening to drink beer tomorrow.

2.4 Assertability has to do with type

So, what is it that distinguishes unassertable antecedents from assertable antecedents in English? I want to suggest that it is *type*, in line with the proposal in (7), repeated here below as (25):

(25) a. stative predicate: $\langle s, t\rangle$
b. eventive predicate: $\langle s{\times}s, t\rangle$

That is, the reason that true eventives are unassertable in English is a formal reason; specifically, they are unable to take a situation argument, because they are of the wrong type to do so.

One argument for this is provided by the fact that only in the simplest phrases does the unassertable-assertable distinction correlate with a conceptual eventivity-stativity distinction. What I mean by this is that it's easy to say that (true) verb phrases about events in English are unassertable, and verb phrases about states in English are assertable; yet we see very similar behavior in larger phrases which we might be more hesitant to say are about events or states. That is, these larger phrases—aspectual, modal, or tensed—absolutely behave like "eventives" or like "statives" on the eventive-stative tests and on assertability. Thus, we must conclude that if these larger phrases are to be thought of as denoting stative or eventive predicates, it is in a very abstract sense, at some distance from any *conceptual* distinction between what states are and what events are. The more abstract distinction seems well-accounted for by a formal property, such as type, as I am proposing here.

Let's take a brief look at some of these larger phrases. With either progressive or perfect aspect added, we get stativity, and also assertability, as in (26). Yet there has been some head-scratching over the nature of the states introduced by aspects (e.g. Hallman 2009); they pass tests for stativity, but what conceptual state do they represent?

(26) a. Xanthippe is writing a letter.
b. Xanthippe must be writing a letter.
c. If you care to know, Xanthippe is writing a letter.

(27) a. Xanthippe has been writing a letter.
b. Xanthippe must have been writing a letter.
c. If you care to know, Xanthippe must have been writing a letter.

Futures provide another example of how larger phrases can have assertability distinctions.[8]

(28) a. #Oh look—it'll rain.
b. Oh look—it's going to rain.

(29) a. #If you care to know, it'll rain.
b. If you care to know, it's going to rain.

The similarity to aspect is such that Copley (2009) analyzed it in terms of aspectualized and unaspectualized futures; as Klecha (2011) points out, however, necessary vs. optional modal subordination is another way to think about it. Either way, if there is a state involved at the highest level of *be going to*, again, it is quite an abstract one.

Past tense poses an interesting case. With past tense in English we only get assertability, and the "stativity" tests are passed handily (*have* here expresses anteriority (Condoravdi 2002)).

(30) a. Xanthippe left.
b. Xanthippe must have left. — epistemic reading possible
c. If you care to know, Xanthippe must have left.

Yet, no one would say that past tense sentences are stative. And actually, discourse facts indicate clearly that they are not, since they advance the narrative time while statives do not (ter Meulen 1995: a.o.), as shown in (31). In (31a), Xanthippe leaves after the time the speaker walks in, while in (31b) Xanthippe is there when the speaker walks in.

(31) a. I walked in. Xanthippe left.
b. I walked in. Xanthippe was there.

So we might say that past tense is transparent to the eventivity/stativity distinction; it just passes up the eventivity/stativity of its complement.

The data for combinations of aspect, modals, futures, and tense are of course complex; we have only scratched the surface. But to recap what I want to take away from this, which we can see already in the few examples above: Assertability and non-assertability correlate

8 The epistemic *must* test is not available with futures, presumably for morphological reasons, and the *oh look* in (28) is required to rule out an assertable reading of *will* that relies on evidence of long standing; see Copley (2009) for details.

nicely with eventivity and stativity with the simplest (verb) phrases, but as more material is added, the assertability/non-assertability distinction remains, while any relevant notion of "stative" or "eventive" gets more and more abstract.

This general effect looks something like a grammatical bleaching effect (Meillet 1912, Traugott 1980, Sweetser 1988). The way I would like to think about it here is that the cognitive ontology and the grammatical ontology do not always have to be the same (Borer 2005, Copley & Harley 2015). One way to demonstrate this point is from mismatches between the two. For instance, grammatical gender and conceptual gender do not always match up (e.g. German *das Mädchen* 'DET$_{neuter}$ girl'), or grammatical mass/count vs. conceptual mass/count (e.g. *There's dog all over the couch.*). Typically, the conceptual categories are rich, while the grammatical categories are discrete and narrow.[9]

What I want to suggest here is such a split between grammatical and conceptual categories, in the service of arguing that type is responsible for the eventive-stative difference, in terms of the eventive-stative tests. The grammatical types "state" (for us, s, a situation) and "event" ($s \times s$ is what we are aiming at) are presumably what the tests diagnose. These are indeed associated, in the verb phrase, with rich conceptual "states" and "events"; but the farther we get from the verb, the less the rich concepts of "state" and "event" matter to "stativity" and "eventivity" in terms of the eventive-stative tests and assertability.

If we now consider, very briefly, what assertion might be, we can see that the notion of type is a good candidate for the kind of formal property required by assertion.

Intuitively, asserting requires the speaker to present a proposition in a certain way. The proposition itself is something that can have a truth value, and depending on one's ontology, this can be achieved in a number of ways. A popular way to do this, especially in linguistic semantics, is to treat propositions as sets of possibilia such as worlds or situations. In this view, propositions are essentially cases of predication (or equivalently, "description"): they partition a domain of possibilia into a set and its complement, corresponding respectively to the truth values true and false.

There is broad consensus, however, that assertability requires more than a mere predication; it additionally requires an act of something like endorsement or commitment from a mind. Conversely, it is quite possible to have a predication without this act. This point can be seen perhaps most clearly in work by Recanati (e.g. Recanati 2007), who argues that assertion requires the speaker to apply the proposition—itself a predicate of situations—to the current situation. We can call this act a "Recanati reckoning".

On this hypothesis, what goes wrong with our unassertable eventives is that they are not predicates of situations, i.e. they are not type $\langle s,t \rangle$, so assertion, which applies a predicate of situations to the current situation, is not available. Consider the word *ball*, the denotation of which, let's assume, is a predicate of entities, type $\langle e,t \rangle$. The utterance of *ball* on its own is not an assertion. It can certainly be uttered, and there are certainly conditions under which its utterance is more or less pertinent, such as when the speaker is urgently pointing at a ball—which suggests that there are conditions on the *utterance* of the word *ball* by itself—but these utterance conditions cannot be assertability conditions in the sense of a Recanati reckoning. The idea is that *ball* is of the wrong type to be asserted.

9 Both, however, can be explicated using formal techniques.

I want to suggest that English unassertable eventives run along the same lines. This would require eventives to be predicates of something other than situations, while statives would have to be predicates of situations (because they are assertable). The prime candidate for this other kind of object is that they are "events", and thus that these eventives are indeed predicates of "events", just as it says on the tin. But how do we know what the type of events is?

3 True eventives are relational

We have now seen evidence that English true eventives have a different type from English statives. English true eventives cannot be asserted, which, I argued above, suggests that they are not predicates of situations. We must look at a different kind of argument to try to understand what the type of eventives actually is.

The usual reason given in the aspectual literature (e.g., Comrie 1976, Giorgi & Pianesi 1997, Smith 1991) for the unassertability of true eventives in English is that there is a problem with having an eventive non-aspectualized predicate occur at the present instant because eventive predicates represent change and it is not possible to represent change at one instant. In other words, eventive predicates *at some level at least*, whether the grammatical level, the conceptual level, or both, should involve a relation, change or comparison between two things; while stative predicates should make reference to only one thing. Such a claim is made succinctly by Hallman (2009), who gives a formal (i.e., grammatical) analysis of this kind.

Many mainstream formal Davidsonian approaches adopt something like this principle to explain the unassertability of English eventives, but the principle is not itself formalized into what the grammatical system sees. That is, there is nothing representing a relation, a comparison, or a change in an expression of the form *predicate*(e), though there certainly can be in the concept of "event" that e refers to. Telic predicates may have subevents, but for atelic eventives, generally there aren't two things but only one thing, namely an event e, and thus the expression *predicate*(e) is insufficient on its own to explain why atelic eventives are also impossible in English at the present moment, i.e. *Xanthippe plays baseball* can't mean that Xanthippe is playing baseball. If we want the principle to fall out of English's grammatical system—and we definitely do, since not all languages are like English—it is necessary that our type difference between eventives and statives make eventives involve two of something and statives involve one of something.

Let's call eventives that are predicates of a relation "relational eventives". Despite the popularity of Davidsonian event arguments, where the eventive is ultimately a predicate of an event argument, the idea of relational eventives has been proposed a number of times, rather independently each time. The (at least) two indices necessary for a relational eventive are represented either as temporal variables (which includes Hallman's interval-based analysis of eventives) or as situations. Theories where the indices are conceptually treated as times include, for instance, Croft (2012) and Verkuyl (2019). Theories where the two indices are treated more like situations include Fernando (2004, 2005), van Lambalgen & Hamm (2008) and Copley & Harley (2015, 2021). (See also Krifka (2014) for similar indices used for certain conditionals.)

For our event semantics, we will use a simplified version of what is proposed in Copley & Harley (2015, 2021). In these, event arguments are replaced with function terms—very dumb functions from one situation to another. Again making use of the grammatical-conceptual distinction, the (grammatical) function is associated with a (conceptual) input of energy that may or may not provoke a change from the initial situation to the final situation. Thus here the "two things" are two situations related causally.

In this paper, we will not represent the input of energy itself. That is, we will treat events as mere ordered pairs. The reason we *can* make events in our toy semantics ordered pairs is because we are dealing with perhaps the simplest eventive predicate of all, namely *become* P. While many predicates place requirements on the manner in which we get from one situation to another, *become* P is one of the ones that does not; in particular, it only says that P does not hold in the first situation and holds in the second. Thus, an ordered pair is sufficient to account for this requirement in our toy semantics here.

With this framework in mind, we can return to the difference between stative predicates (type $\langle s, t\rangle$) and eventive predicates (type $\langle s{\times}s, t\rangle$, as well as understanding how this fits into what we know about assertion.

True eventives in English can't be asserted, I have argued, because they are the wrong type; while they are of type $\langle s{\times}s, t\rangle$, assertion is looking for something of type $\langle s, t\rangle$. But even though assertion is not possible for eventives we can imagine a different way to use a type $\langle s{\times}s, t\rangle$ predicate to update the situation—in effect, one is not adding a proposition, as in assertion, but noting that an event occurs (or has occurred, or is occurring). So, formally, assertion and this kind of noting (I will call it "happening") will share something even though they are different. We will update in two ways—one way with type $\langle s, t\rangle$ (i.e. adding a proposition) and one way with $\langle s{\times}s, t\rangle$.[10]

4 Solving the puzzle

We are ready to solve the puzzle. The idea is that the causal relation inherent to true eventives is, literally, the causal relation from which springs the causal readings of conditionals like that in (1b). The proposal involves a mapping from a causal model to the event semantics of Copley & Harley (2021). We will do this by indexing the nodes in the causal model and similarly indexing a sequence of situations that will easily map to a relational meaning for eventives. Finally, to show how this mapping between linear causal models and event semantics resolves the puzzle in (1), we will use a very simple dynamic semantic denotation for conditionals; more sophisticated dynamic treatments such as Kaufmann (2000) are not necessary here.

Incidentally, the indexing is why, in this paper, we can only use a "toy", non-branching (linear) causal model. We are greatly aided by the fact that our causal model has causality going in a single direction ($C \rightarrow Q \rightarrow P$); in other words, we don't have to deal with the

10 Usually a distinction is made, in theory as in language, between an entity's information state about the world, and the state of the world itself. The system I will be presenting here seems to do its job without such a distinction; the situations here represent only the way the world is. However, inferences about these situations can be made by using information from the causal model, without reifying information states. I am not sure yet whether this setup amounts to a feature or a bug; the question turns on whether it is necessary or otherwise helpful to reify information states *for these particular examples.*

complications of a branching network of causal relations with "colliders" ($A \rightarrow B \leftarrow C$) or "forks" ($A \leftarrow B \rightarrow C$). It's simple to inductively define a sequence of situations from the sequence of nodes; harder to say how to index situations from a network of nodes. Do we need to use differently-defined situations? Times? And moreover, the answers to these questions may only come, or may come most interestingly, from cross-linguistic data. Thus, in this paper we will stick to a linear causal model. However, the definitions below are in principle useable for branching models, if only the indexing definition can be appropriately changed.

4.1 Causal models with indexed situations

The next two definitions are simplified, sometimes radically, from Pearl (2000: p. 44) and inspired by similar adaptations from Baglini & Francez (2016), Baglini & Bar-Asher Siegal (2019).

Definition 1 A *causal structure* $\mathcal{D}$ *of a set* $\mathcal{V}$ *of variables* is a directed acyclic graph (DAG) in which each node corresponds to a distinct element of $\mathcal{V}$. Links between the nodes are represented as ordered pairs of elements of $\mathcal{V}$. Each link represents a direct functional relationship among the corresponding variables.

The usual kinship terms (notably "parent" and "child") can be used on these structures. We can also define (in a simplified fashion) *endogenous* and *exogenous* variables:

Definition 2 In a causal structure of a set $\mathcal{V}$ of variables, a variable A is *endogenous* if there is a variable X such that (X, A) is a link. Conversely, A is *exogenous* if there is no variable X such that (X, A) is a link.

Causal models are to be defined as follows:[11]

Definition 3 A *causal model* is a pair $\mathcal{M} = \langle \mathcal{D}, \Theta_{\mathcal{D}} \rangle$ consisting of a causal structure $\mathcal{D}$ and a set of parameters $\Theta_{\mathcal{D}}$ compatible with $\mathcal{D}$. The parameters in $\Theta_{\mathcal{D}}$ assign a function $d_A = f_A(\text{parents}_A)$ to determine the value of each node V, where parents$_V$ is the set of parents of A in $\mathcal{D}$.

The idea behind Definition 3 is simply that by definition, in a causal model, the value of any node A depends on the value of its parent(s); in the example we are dealing with, of course, Q and P each only have one parent and C is endogenous, so has no parent. We will skip over the question of what exactly the set $\Theta_{\mathcal{D}}$ of parameters corresponds to. It can be thought of as a placeholder for whatever it is that determines the value of the child given the value of the parent.

11 Definition 3 is really very simplified indeed with respect to structural causal models. Usually one would include two other factors: an "error" variable such that the child depends on *both* the set of parents and the error variable, as well as a condition that when conditioned on the parents, the child is independent of its non-descendants (the "causal Markov condition"). The latter is a constraint on which models are useful for calculating causal influence; those that do not obey this condition do not seem to be appropriate for causal reasoning. Because the argmentation for this is involved, we will elide it here, but see Pearl (2000) for extensive discussion.

The relations in the causal model behave as in Pearl (2000), where a relation between nodes such as (A, B) conveys that the second node "listens" to the first node. That is, the second node's value is sensitive to first node's value. It is very important to remember that this relation is not always paraphraseable by the main verb *cause*; better words are *influence* or *affect*. The absence of such a relation between nodes conveys that there is no influence from one node to another.

Now, for our toy model, we can further winnow down causal models to those that are linear. The model represented graphically above in (2) is a linear causal model.

Definition 4 A *linear causal model* is a causal model such that for all its links with arbitrary A, B such that (A, B), and for arbitrary X: if $(X, B) \in$ model, then $A = X$, and if $(A, X) \in$ the model, $X = B$.

Taking a step back, let's remind ourselves what we need causal models to do. Causal models need to constrain the set of possible situations that we are considering, that can be the case. That is, they need to give us at least the information in the table we saw above:

(32)

	C	Q	P
a.	1	1	1
b.	0	0	0
c.	1	0	0
d.	1	1	0

Causal models have another job to do, however. They also need to constrain possible *transitions* between situations in a causal sequence, i.e., possible (relational) events. This is going to be the key to linking causal models to event semantics. In doing this, they will also give us the information in (32)—that is, the information in the table in (32) above will be generated by what we are about to do in (33) below. First we will define our fundamental types and some expressions having to do with predication.

Types. We will use several types: entities (type e, variables $x, y, \ldots$); situations (type s, variables $s, s', \ldots$); and type t which will represent truth values. The situations are, intuitively, to be identified with lines on the truth table, in other words, partial assignments of values to basic predications. The degree type will correspond to truth values for our toy semantics in this paper. We will use the type e to represent the complex type $s{\times}s$ which represents events, with variables $e, e', \ldots$.

Basic predications. For the nodes of our causal model, we will first distinguish between what we will call *basic predications* (cf. *measure functions* (cf. e.g. Hay et al. (1999))), and *relations*. Basic predications are of the form $p(x)(s)$, where p is a function from entities and situations to degrees (values). A *basic predication equation* gives the value of this function evaluated at these arguments. The word "basic" emphasizes that there is no predication more basic than that which is evaluated at a situation; that is, we will not be writing anything like "$p(x)$".

Definition 5 For any p, x, s where p is a function from entities and situations to degrees, x is an entity, and s is a situation, an expression of the form $p(x)(s)$ is a *basic predication expression*.

Definition 6 An equation of the form $p(x)(s) = d$ is a *basic predication equation.*

Definition 7 An expression of the form $\lambda s.p(x)(s)$ is a *basic predication lambda expression*; it names a function that takes a situation s and returns the degree to which x is p in s.

Nodes map to basic predication lambda expressions. Here we get into the crucial part that allows us to link causal models with event semantics. The nodes in our causal model are going to be interpreted, via an assignment function $\mathcal{G}$, as basic predication lambda expressions, as in Definition 7. So, where we are used to seeing a simple variable with a value, i.e., something like $X = x$, we have now "replaced" the variable with a basic predication lambda expression.

It may not be clear yet why we would want to put situation arguments into the interpretations of nodes. The reason why stems from our working assumption that eventives are predicates of relations from situations to situations; these are the situations at which basic predications are evaluated. So, since we know that eventive verbs are in a sense instances of the arrows in a causal model but with the values of the nodes filled in, and in causal models the arrows are relations between nodes, it makes sense that we need some way to get from nodes to situations. The assignment function $\mathcal{G}$ will do this job for us. It assigns to every node a corresponding expression of type $\langle s, t \rangle$.

(33) $\quad \mathcal{G}(A) = \lambda s.p(x)(s)$

Now our linear causal model needs to build a situation sequence for us. We can use the directions of the arrows to construct by induction a set of basic predication expressions, using the basic predication expressions from each node but indexing situations in each equation. In particular, indices $n, n+1, \ldots$ are added along the direction of the arrows. Then we will assign the same indices to the variables. The indices themselves have no significance except to provide a successor.

Definition 8 An indexing on a *linear causal model* $\mathcal{M} = \langle \mathcal{D}, \Theta_{\mathcal{D}} \rangle$ is defined inductively as follows:

An exogenous variable is assigned the index n (write: A_n). For all B is such that (A, B), B is assigned the index $n+1$ (write: B_{n+1}). With such an indexing, $\mathcal{M}$ is an *indexed linear causal model.*

Definition 9 To define a *causal set of basic predication expressions* for $\mathcal{M}$, let $\mathcal{V}$ be a set of variables; $\mathcal{M}$ be an indexed linear causal model for $\mathcal{V}$; and $\mathcal{G}$ be an interpretation function assigning each variable in $\mathcal{P}$ a basic predication lambda expression. Then, for any node A_n such that A_n is exogenous in the indexed linear causal model M, add $[\mathcal{G}(A)](s_n)$ to the set. For any node B_{n+1} such that B is endogenous in the indexed linear causal model M and add $[\mathcal{G}(B)](s_{n+1})$ to the set.

(34) $\quad$ Causal set of basic predication expressions $\{c(h)(s_n), q(y)(s_{n+1}), p(x)(s_{n+2})\}$

Along with the few assumptions we made about Clavarelian dynastic politics that gave us the model in the first place, there are some further common-sense assumptions to make.

(35) a. Persistence: Once a queen/princess, always a queen/princess
b. Closed World: If not mentioned, then 0

(35a) is a fact about what it's like to be a queen or a princess in Clavarel. (35b) is a more general constraint. We will need to apply them in the order given (first Persistence, then Closed World) in our derivations to get the right result.

Persistance gives us a different set of basic predications:

(36) Entailed set of basic predications for persistence
$\{q(y)(s_{n+2}), p(x)(s_{n+3})\ldots\}$

4.2 Denotations of verbal predicates

Now we can turn to the language side of things, beginning with verbal predicates.

We will assume that the meaning of a stative predicate such as *be the princess* is essentially a basic predication.

(37) $[\![\text{be the princess}]\!] = \lambda x \lambda s. [\![\text{the princess}]\!](x)(s)$

Such a predicate takes an entity as its argument to yield something of type $\langle s, t \rangle$, i.e., a basic predication.[12]

Since events are ordered pairs of situations, eventive predicates are predicates of ordered pairs of situations. Let us consider an arbitrary ordered pair $e = (1_e, 2_e)$ where 1_e and 2_e are both situations. Then *become the princess* has the following denotation, keeping in mind that this is interpreted as involving energetic causation, although we are only reifying the change, not the energy:

(38) $[\![\text{become the princess}]\!] = \lambda x \lambda e \in D_{s \times s}. \neg [\![\text{the princess}]\!](x)(1_e)\ \&\ [\![\text{the princess}]\!](x)(2_e)$

Again, I stress that the relational argument of eventive predicates here is not going to correspond to one of the arrows in a causal model, but rather to a relation between situations chosen from a sequence of situations that is constrained by that causal model. Causal models provide the structures that are further constrained by eventive predicates. Recall too that the nodes in causal models do not need values in order for that formal structure to be a causal model. To reiterate: **events are not to be identified with the arrows in causal models**, even though they are similarly represented as ordered pairs. While a causal model provides information as to the possible situations and ordered pairs of situations that are compatible with it, it does not pick out a single ordered pair of those situations; this is what an event is. So rather than corresponding to arrows in the causal diagram, events

12 For readers unfamiliar with compositional semantic notation (and see Heim & Kratzer (1998) for an excellent introduction), the evaluation function $[\![\,]\!]$ relates linguistic expressions to their meanings (or "denotations"). In lambda (λ) expressions as in the right half of the equation in (38), the lambdas allow us to correctly compose meanings together according to the syntactic structure they are in. So, in (38), the meaning of *be the princess* first takes an entity argument x (in our example, x refers to Xanthippe) and then takes a situation argument s. This order of operations is also reflected in the syntactic structure, where compositional operations are only possible between "sisters" in the tree. It is possible to leave the lambdas out here and still understand the proposed solution to the puzzle, but semanticists very much wish to see them there to understand how to compose the meanings.

correspond rather to transitions between situations, where the set of possible transitions between situations at any point is constrained by the causal model.

The interpretation of eventive predicates as involving energy tells us that we will be able to say that we can move between situations (update them) through the input of energy. But this is not the only way to move between situations.

Events and assertions thus both take us from one situation to another, but with vastly different interpretations of how we get from one situation to the other: either through energetic causation, which happens in the world, and which puts a condition on the initial situation; or through a mental process of concluding/learning a proposition, which does not put such a condition on the initial situation.

4.3 *If*

For *if*, we need a meaning that will meet two conditions. First, it needs to allow each prejacent to update in its own way (learning or happening). This can be accomplished through compositionality. Second, it needs to ensure that any possible "routes" where φ survives but ϕ does not are removed.

Our toy causal model, however, is linear, and so is the situation sequence associated with it. The situation sequence gives us a single name for the immediately causally accessible situation from s_n – there is only one, namely s_{n+1}. Where $A \to B$ is in the model, then provided that we know whether A holds at s_n, we know whether B holds at s_{n+1}, because that is what the causal model tells us. Effectively, the work done by the universal quantifier in (38) is here done by the causal model. Here, we don't need to use a quantifier to remove from s any possibilities where φ survives but φ does not, because we already know from the causal model that if φ survives, there is no possible outcome where ψ does not survive. (If alternatively we were to indeed consider outcomes where φ survives yet ψ does not, we would need to assume a different causal model from the one we are assuming.)

Thus, all we *need* for the toy system is the conjunctive update in (39) and this is the update I will use.

(39) $s[\text{if } \varphi, \psi]^{\mathcal{M}} := s[\varphi]^{\mathcal{M}}[\psi]^{\mathcal{M}}$

A situation is updated with φ and then the result of that is updated with ψ to yield another situation. All of this takes place with respect to the causal model M. As usual, a successful update is one where the update results in a situation in the model and this yields truth (or if one prefers, acceptability). Note that s must be in the domain of φ as we would expect with indicative conditionals.

This move does however raise some questions.

First, it raises the question of whether even a non-linear causal model would remove the routes where φ survives and ψ does not. The answer here depends on whether knowing the value of an influencing node A determines the value of the influenced node B, where $A \to B$ is in the model. This in turn depends on whether the arrows correspond to functions. Typically this is not taken to be the case; the value of a node is a function of the values of all its influencing nodes. However, models that do have this property have been explored (Copley & Kagan 2021).

Second, even if the conjunctive update were to be generally appropriate, it would immediately raise the question of how to express the difference between conjunction and implication. I can point to the limited use of conjunctive structures having conditional meaning as a suggestive fact linking conjunctions to implications, but it is merely a suggestive fact at this point. I do think that an answer to that question could have something to do with *if* encoding reference to a causal model and *and* not encoding it, but I will not try to make this thought more precise here.

Moving on, we can further define a relation for compatibility/accessibility. This will be used in the "learning" kind of update below. This definition ensures that one can learn that P (I assume that truth values do not change once learned).

(40) compatibility/accessibility relation R: for arbitrary situations s, s', sRs' just in case all basic predications mentioning s of the form $p(x)(s)$ are such that the truth value of $p(x)(s)$ is equal to the truth value of $p(x)(s')$.

Our two ways to update, one way for statives and one way for eventives, are in (41) below.[13]

(41) Two kinds of update:
a. learning: $s[\varphi_{\langle s,\, t\rangle}]^{\mathcal{M}} := \iota s' : sRs' \text{ and } \varphi(s')$
b. happening: $s[\varphi_{\langle s\times s,\, t\rangle}]^{\mathcal{M}} := \iota s' \in \mathcal{S}^{\mathcal{M}} : [\iota e \in \mathcal{S}^{\mathcal{M}}\times\mathcal{S}^{\mathcal{M}} : [2_e = s' \text{ and } \varphi(e)]]$

Thus, depending on whether φ is a stative (type $\langle s, t\rangle$, (41a)) or an eventive (type $\langle s\times s, t\rangle$, (41b)), different updates are used. Note that e, which represents the event, and is of type $\langle s\times s, t\rangle$, is in a very real sense analogous to the accessibility/compatibility relation in (41a). Another important point is that both of these put a constraint on the starting situation s. In (41a), s has to be in the domain of R, and in (41b), s has to be in the domain of e. However, in practice, the accessibility/compatibility relation is more permissive than *become*. What this corresponds to is the fact that with *being* in s', s' must only be accessible to/compatible with s. So, in s the value of φ is either true or its value is unknown. However, *become the princess*, as in (38), requires that in the (causally) previous situation s, φ is false. This point turns out to be crucial to solving the puzzle.

4.4 Putting it all together

Recall that we want to explain (1a) and (1b), repeated here as (42a) (which is true) and (42b) (which is false):

(42) a. If Xanthippe is the princess, Yolanda is the queen.
b. If Xanthippe becomes the princess, Yolanda becomes the queen.

Informally, in both, we want the antecedent to take us from the current situation to another one, and the consequent takes us from that other situation to yet another one; we update first with the antecedent φ, then with the consequent ψ.

13 There are in principle a couple of alternative ways to achieve the same effect; we could have instead put it in a dedicated assertion operator, or in a verbalizing head.

The procedure for evaluation is to search the causal direction set of basic predication expressions given by the causal model for the terms needed. This set, whose elements we last saw populating the table in (34), is represented again in (43):

(43) Basic predication expressions, causal direction ($\mathcal{M}$)

$$\begin{aligned} s_n: &\quad c(h)(s_n) \\ s_{n+1}: &\quad q(y)(s_{n+1}) \\ s_{n+2}: &\quad p(x)(s_{n+2}) \end{aligned}$$

So, in (1a), to evaluate the antecedent we search for something in M that tells us the value of $p(x)(s)$ for any s. We have $p(x)(s_{n+2})$, so we start at s_{n+1}. Updating that situation with the antecedent returns s_{n+2}. Now we check to see if we can update s_{n+2} with something that tells us the value of $q(y)(s_{n+2})$. In fact, because of our assumption of persistence (once the queen always the queen), we can learn that $q(y)(s_{n+2})$ is true.

Let $\varphi = p(x)$ and $\psi = q(y)$.

(44)
a. $s[\text{if } \varphi_{\langle s,t\rangle}, \psi_{\langle s,t\rangle}]^{\mathcal{M}} = s[\varphi_{\langle s,t\rangle}]^{\mathcal{M}}[\psi_{\langle s,t\rangle}]^{\mathcal{M}}$ — definition of *if*
b. $s[\varphi_{\langle s,t\rangle}] = \iota s' \in S^{\mathcal{M}} : sRs' \text{ and } \varphi_{\langle s,t\rangle}(s')$ — definition of learning
c. $\iota s' \in S^{\mathcal{M}} : [sRs' \text{ and } \varphi_{\langle s,t\rangle}(s')] = s_{n+2}$ — since $\varphi_{\langle s,t\rangle} = p(x)$ and $p(x)(s_{n+2})$
d. $s[\text{if } \varphi_{\langle s,t\rangle}, \psi_{\langle s,t\rangle}]^{\mathcal{M}} = s_{n+2}[\psi_{\langle s,t\rangle}]^{\mathcal{M}}$ — from (44a) and (44c)
e. $s_{n+2}[\psi_{\langle s,t\rangle}] = \iota s' \in \mathcal{S}^{\mathcal{M}} : s_{n+2}Rs' \text{ and } \psi_{\langle s,t\rangle}(s')$ — definition of learning
f. $\iota s'' \in \mathcal{S}^{\mathcal{M}} : [s_{n+2}Rs'' \text{ and } \psi_{\langle s,t\rangle}(s'')] = s_{n+2}$ — since $q(y)(s_{n+1})$ and persistence

The result of these two updates is s_{n+2}, which is in $\mathcal{S}^{\mathcal{M}}$, so this utterance is judged true/acceptable, as desired.

For (1b), unlike (1a), we have to have it that Xanthippe is not princess in the starting situation. So we find ourselves starting in situation s_{n+1}. Then the antecedent takes us to s_{n+2}. Can we say anything about our consequent evaluated at that situation, namely $q(y)(s_{n+2})$? Due to persistence, we can; it is true. This makes it true that Yolanda is the queen in s_{n+2}, but this is not what we were looking for. We were looking for Yolanda to *become* the queen. For this to be true, we would need to know that $q(y)(s_{n+2})$ were *false*, and that $q(y)(s_{n+3})$ were true. This is not the case so (1a) is false. And since we know that $q(y)(s_{n+1})$ is false and $q(y)(s_{n+2})$—that is, that Yolanda becomes queen from s_{n+1} to s_{n+2}—we can see where the impression that (1b) is exactly causally backwards comes from.

(45)
a. $s[\text{if } \varphi_{\langle s\times s,t\rangle}, \psi_{\langle s\times s,t\rangle}]^{\mathcal{M}} = s[\varphi_{\langle s\times s,t\rangle}]^{\mathcal{M}}[\psi_{\langle s\times s,t\rangle}]^{\mathcal{M}}$ — definition of *if*
b. $s[\varphi_{\langle s\times s,t\rangle}]^{\mathcal{M}} :=$
$\iota s' \in \mathcal{S}^{\mathcal{M}} : [\iota e \in \mathcal{S}^{\mathcal{M}}\times\mathcal{S}^{\mathcal{M}} : 1_e = s \,\&\, 2_e = s' \,\&\, \varphi(e)]]$ — def. happening
c. $\iota s' \in \mathcal{S}^{\mathcal{M}} : [\iota e \in \mathcal{S}^{\mathcal{M}} \times \mathcal{S}^{\mathcal{M}} : [1_e = s \,\&\, 2_e = s' \,\&\, \varphi(e)]] = s_{n+2}$
d. $s[\text{if } \varphi_{\langle s,t\rangle}, \psi_{\langle s\times s,t\rangle}]^{\mathcal{M}} = s_{n+2}[\psi_{\langle s\times s,t\rangle}]^{\mathcal{M}}$ — from (45a) and (45c)
e. $s_{n+2}[\psi_{\langle s\times s,t\rangle}]^{\mathcal{M}} =$
$\iota s'' \in \mathcal{S}^{\mathcal{M}} : [\iota e' \in \mathcal{S}^{\mathcal{M}} \times \mathcal{S}^{\mathcal{M}} : [1'_e = (s_{n+2}) \,\&\, 2'_e = s'' \text{ and } \psi(e')]]$ def. happening
f. $\iota s'' \in \mathcal{S}^{\mathcal{M}} : [\iota e' \in \mathcal{S}^{\mathcal{M}} \times \mathcal{S}^{\mathcal{M}} : [1'_e = (s_{n+2}) \,\&\, 2'_e = s'' \,\&\, \psi(e')]] = s_{n+3}$

The result of these two updates is s_{n+3}. But even if we accommodate the existence of an s_{n+3} (we need to redefine $\mathcal{S}^{\mathcal{M}}$ to allow for this but this should not be problematic), we still have a problem. The problem is the contradiction between (46a) and (46b):

(46) a. Yolanda is not queen in s_{n+2} — because s_{n+2} must be in in domain of f'
b. Yolanda is queen in s_{n+2} — because of persistence

So, since the consequent update we are trying to make results in contradictory demands on s_{n+2}, the conditional is judged false/unacceptable. The toy semantics here thus explains both (1a) and (1b) and thereby solves our puzzle.

4.5 Related examples

There is not too much more we can do with the linearity restriction on our causal models. Classic tests such as strengthening the antecedent and Sobol sequences, for instance, will require non-linear causal models, which we cannot yet link to event semantics. However, there are a few related examples that we can more or less easily treat with a linear causal model and which are likely to be questions in the mind of the reader at this point.

One question is what happens if we mix and match the eventives and statives as in (47a) and (47b).

(47) a. If Xanthippe becomes the princess, Yolanda is the queen. — true
b. If Xanthippe is the princess, Yolanda becomes the queen. — false

The framework correctly accounts for these judgments. In (47a), Xanthippe becomes the princess from s_{n+1} to s_{n+2}, and by persistence, Yolanda is indeed the queen in s_{n+2}. For (47b), Xanthippe is the princess in s_{n+2}, but from s_{n+2} we can't do a happening update with Yolanda's becoming the queen, for the same reasons as for (1b).

A second question picks up some loose ends from above. How does this framework deal with the futurate, generic/habitual, and storytelling readings, which I argued above are really stative? Given that they are stative, they are predicates of situations and therefore they update via learning. So, for instance, the example in **??** above, repeated here in (48), gets the reading that on updating via learning that Xanthippe has a plan to leave tomorrow, we update via learning that Zelinda calls her today.

(48) If Xanthippe leaves tomorrow, Zelinda calls Yolanda today.

A third question that we can answer here has to do with past tense; what happens if we slightly complicate our simplest conditionals by adding past tense to one or more of the prejacents? We have avoided time until now (and correctly, I think, for the examples we were dealing with) but we cannot avoid it forever. So, without getting into questions of temporal anaphora which might make some of the mixed and matched conditionals with past tense feel infelicitous, here is a felicitous conditional which is true relative to our causal model M.

(49) If Xanthippe became princess, Yolanda became queen.

Recall that matrix sentences with past tense on them are generally assertable. That means that both prejacents here are predicates of situations, in our toy semantics; the result in (50) is the same (using here an existential quantifier theory of tense, as in (50)). The assertability of the past tense phrases tells us we need to look at a single situation, not a relation between situations, to evaluate them. They are predicates of situations, not predicates of ordered pairs.

(50) $\textsc{Past}(\mathrm{p}) = \lambda s.\exists s'\ \textsc{before}\ s : p(s')$

Consequently, for (50), both prejacents must involve a learning update, rather than a happening update. The conditional has us first learning that, in a past situation relative to s, Xanthippe became princess, so we update s with that fact; consequently we learn that in a past situation relative to the updated situation, Yolanda became queen. This is congruent with our model, so the sentence is judged true.

Likewise, the framework can also handle "backtracking" conditionals such as (51). The antecedent, which has a true eventive prejacent, is updated via happening, while the past tense consequent is updated via learning.

(51) If Xanthippe becomes princess, Yolanda became queen.

So, the predicted reading is that the antecedent takes us from s where Xanthippe is not princess to s' where she is ($= s_{n+3}$) and then we learn, as for (50), that in a past situation relative to the updated situation, Yolanda became queen. This is the case, so (51) is judged true.

References

Austin, John Langshaw. 1961. Ifs and cans. In *Philosophical papers*, 153–180. Oxford University Press.

Baglini, Rebekah & Elitzur A. Bar-Asher Siegal. 2019. Direct causation: A new approach to an old question.

Baglini, Rebekah & Itamar Francez. 2016. The implications of managing. *Journal of Semantics* 33. 541–560.

Bar-Asher Siegal, Elitzur & Nora Boneh. 2019. Sufficient and necessary conditions for a non-unified analysis of causation. In *36th west coast conference on formal linguistics*, 55–60. Cascadilla Proceedings Project.

Bennett, Jonathan. 1988. Farewell to the phlogiston theory of conditionals. *Mind* 97(388). 509–527.

Bennett, Jonathan. 2003. *A philosophical guide to conditionals*. Oxford University Press.

Borer, Hagit. 2005. *Structuring sense: Volume i: In name only (structuring sense)*. Oxford University Press.

Carlson, Greg. 2009. Generics and concepts. In Francis Jeffry Pelletier (ed.), *Kinds, things, and stuff: Mass terms and generics*, 16–35. Oxford University Press.

Comrie, Bernard. 1976. *Aspect*. Cambridge University Press.

Condoravdi, Cleo. 2002. Temporal interpretation of modals. In David Beaver, Stefan Kaufmann, Brady Clark & Luis Casillas (eds.), *Stanford papers on semantics*, Palo Alto: CSLI Publications.

Copley, Bridget. 2008. The plan's the thing: Deconstructing futurate meaning. *Linguistic Inquiry* 39(2). 261–274.

Copley, Bridget. 2009. *The semantics of the future*. Routledge. *Outstanding Dissertations in Linguistics*, Routledge.

Copley, Bridget. 2018. Dispositional causation. *Glossa: A Journal of General Linguistics* 3(1).

Copley, Bridget & Heidi Harley. 2015. A force-theoretic framework for event structure. *Linguistics and Philosophy* 38. 103–158.

Copley, Bridget & Heidi Harley. 2021. What would it take to tame the verbal hydra? CNRS/Paris 8 and University of Arizona ms.

Copley, Bridget & Olga Kagan. 2021. The Russian perfective and maximality of causal influences. Talk presented at the Converging On Causal Ontology Analyses (COCOA) come-as-you-are online workshop, June 2021.

Croft, William. 2012. *Verbs: Aspect and causal structure*. Oxford University Press.

Davidson, D. 1967. The logical form of action sentences. In N. Rescher (ed.), *The logic of decision and action*, University of Pittsburgh Press.

DeRose, Keith & Richard E Grandy. 1999. Conditional assertions and" biscuit" conditionals. *Noûs* 33(3). 405–420.

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

Dudman, Victor H. 1983. Tense and time in english verb clusters of the primary pattern. *Australian Journal of Linguistics* 3(1). 25–44.

Edgington, Dorothy. 1995. On conditionals. *Mind* 104(414). 235–329.

Edgington, Dorothy. 2003. What if? questions about conditionals. *Mind & Language* 18(4). 380–401.

Fernando, T. 2004. A finite-state approach to events in natural language semantics. *Journal of Logic and Computation* 14(1). 79–92.

Fernando, Tim. 2005. Comic relief for anankastic conditionals. In P. Dekker & M. Franke (eds.), *Proceedings of the 15th amsterdam colloquium*, 71–76.

Giorgi, Alessandra & Fabio Pianesi. 1997. *Tense and aspect: From semantics to morphosyntax*. Oxford University Press on Demand.

Groenendijk, Jeroen & Martin Stokhof. 1990. Two theories of dynamic semantics. In *European workshop on logics in artificial intelligence*, 55–64. Springer.

Hallman, Peter. 2009. Instants and intervals in the event/state distinction. Unpublished ms., UCLA.

Hay, J., C. Kennedy & B. Levin. 1999. Scalar structure underlies telicity in degree achievements. In *Proceedings of salt*, vol. 9, 127–144. Citeseer.

Heim, Irene. 1982. *On the semantics of definite and indefinite noun phrases*: UMass dissertation.

Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Blackwell.

Higginbotham, James. 1983. The logic of perceptual reports: An extensional alternative to situation semantics. *The Journal of Philosophy* 100–127.

Higginbotham, James. 1985. On semantics. *Linguistic inquiry* 16(4). 547–593.

Higginbotham, James. 1986. Linguistic theory and Davidson's program in semantics. In E. Lepore (ed.), *Inquiries into truth and interpretation*, 29–48. Oxford: Blackwell.

Kamp, H. 1981. A theory of truth and semantic representation. In T.M.V. Janssen J.A.G. Groenendijk & M.B.J. Stokhof (eds.), *Formal methods in the study of language.*, vol. 135 Mathematical Centre Tracts, 277–322. Amsterdam.

Kaufmann, Stefan. 2000. Dynamic context management. *Formalizing the dynamics of information* 171–188.

Kaufmann, Stefan. 2004. Conditioning against the grain. *Journal of Philosophical Logic* 33(6). 583–606.

Kaufmann, Stefan. 2005. Conditional truth and future reference. *Journal of Semantics* 22(3). 231–280.

Kaufmann, Stefan. 2013. Causal premise semantics. *Cognitive science* 37(6). 1136–1170.

Kaufmann, Stefan, Cleo Condoravdi & Valentina Harizanov. 2006. Formal approaches to modality. In *The expression of modality*, 71–106. De Gruyter Mouton.

Kaufmann, Stefan Heinz. 2002. *Aspects of the meaning and use of conditionals*: dissertation.

Klecha, Peter. 2011. Optional and obligatory modal subordination. In *Proceedings of sinn und bedeutung*, vol. 15, 365–379.

Kratzer, Angelika. 1995. Stage-level and individual-level predicates. *The generic book* 125–175.

Krifka, Manfred. 2014. Embedding illocutionary acts. In *Recursion: Complexity in cognition*, 59–87. Springer.

Krifka, Manfred, Francis Jeffry Pelletier, Gregory Carlson, Alice Ter Meulen, Gennaro Chierchia & Godehard Link. 1995. Genericity: an introduction. In Gregory Carlson & Francis Jeffry Pelletier (eds.), *The generic book*, 1–124.

Lakoff, George. 1971. Presupposition and relative well-formedness. In Leon A. Jacobvits Danny D. Steinberg (ed.), *Semantics: An interdisciplinary reader in philosophy, linguistics, and psychology*, 329–340. Cambridge, UK: Cambridge University Press.

van Lambalgen, Michiel & Fritz Hamm. 2008. *The proper treatment of events*, vol. 6. John Wiley & Sons.

Lauer, Sven & Prerna Nadathur. 2018. Sufficiency causatives. *Manuscript, in revision* .

Maienborn, C. 2005. On the limits of the Davidsonian approach: The case of copula sentences. *Theoretical Linguistics* 31(3). 275–316.

Maienborn, C. 2007. On Davidsonian and Kimian states. In I. Comorovski & K. Von Heusinger (eds.), *Existence: Semantics and syntax*, vol. 84 Springer Science and Business Media, 107–130. Springer.

Meillet, Antoine. 1912. L'évolution des formes grammaticales .

ter Meulen, Alice. 1995. *Representing time in natural language*. Cambridge, MA: MIT Press.

Nadathur, Prerna. 2016. Causal necessity and sufficiency in implicativity. In *Semantics and linguistic theory*, vol. 26, 1002–1021.

Parsons, Terence. 1990. *Events in the semantics of English: A study in subatomic semantics*. Cambridge, Mass.: MIT Press.

Pearl, Judea. 2000. *Causality: Models, reasoning and inference*. Cambridge University Press.

Pearl, Judea & Dana Mackenzie. 2018. *The book of why: The new science of cause and effect*. Basic Books.

Pustejovsky, James. 1995. *The generative lexicon*. MIT Press.

Ramchand, G. 2008. *Verb meaning and the lexicon: A first-phase syntax*. Cambridge Univ Pr.

Recanati, François. 2007. *Perspectival thought: A plea for (moderate) relativism*. Oxford University Press.

Ritter, Elizabeth & Sara Thomas Rosen. 1997. The function of *have*. *Lingua* 101(3). 295–321.

Smith, Carlota. 1991. *The parameter of aspect*. Dordrecht: Kluwer.

Sweetser, Eve E. 1988. Grammaticalization and semantic bleaching. In *Annual meeting of the berkeley linguistics society*, vol. 14, 389–405.

Traugott, Elizabeth Closs. 1980. Meaning-change in the development of grammatical markers. *Language sciences* 2(1). 44–61.

Verkuyl, Henk J. 2019. Event structure without naïve physics. In Robert Truswell (ed.), *The oxford handbook of event structure*, 171–204. Oxford University Press.

Vetter, David C. 1973. Someone solves this problem tomorrow. *Linguistic Inquiry* 4:1. 104–108.

Woodward, James. 2006. Sensitive and insensitive causation. *The Philosophical Review* 115(1). 1–50.