

COUNTERDIRECTIONALITY IN THE GRAMMAR

Jyoti Iyer *University of Massachusetts*

BACKGROUND Adverbs like English *back* involve a presupposition that there exist a prior event that is in some way the *reverse* of the asserted event. In *Bilbo went to Rivendell and later came back home*, two events are overtly present, but that need not be the case. *Bilbo went back to the Shire* is only true if Bilbo had gone to the Shire in the past, whether that is overtly mentioned or not. The presupposition projects through negation-and question-holes: *Bilbo did not go back to the Shire*, and *Will Bilbo go back to the Shire?* The denotation below for the abstract adverb counter is what Patel-Grosz & Beck (2014, and *forthcoming*) assume for Kutchi Gujarati *pacho*, same as Fabricius-Hansen (2001) for German *wieder*. (Presupposition is in box.)

(1) RELATION BETWEEN PROPERTIES OF EVENTS

$$[\![pacho]\!] = \lambda P_{st} \lambda e_s. \boxed{\exists e' [\tau(e') \prec \tau(e) \wedge P_C(e')]} .P(e) \text{ where } P_C \text{ is the reverse of } P$$

I argue that counterdirectionality must be understood as a relation between events, because using a relation between properties of events (1) makes some predictions that are not met. **PROBLEM** The main problem is that having an existential definition (presuppose that there exist a property...) tells us nothing about the cases where even though the existential condition is met, the counterdirectional adverb is still bad. In Patel-Grosz & Beck (P-G&B), one type of P and P_C is the situation where e and e' are reverse because we understand them as having swapped Source/Goal of their traversed paths (where Source and Goal can be overt, but need not be).

(Examples below are intended as event sequences, they could have been presupposed, not necessary to be overt.)

Example (2) is fine on this rationale, but if we go to example (3), *back* no longer makes any sense. We are not able to understand Gandalf coming out of the kitchen as the reversal of Bilbo going into the kitchen, even though the Source/Goal are reversed. We expect identity of other properties of the event, but here $\text{Agent}(e) \neq \text{Agent}(e')$. (There is a good but irrelevant reading with accommodation of the presupposition of *back* i.e. a prior event of Gandalf going in to the kitchen. But in this sentence it is not licensed, ∴ has the flavour of a non-sequitur.)

- (2) *ok* Bilbo went to the Shire, then came back. (ruled in by PG&B) ✓
- (3) # Bilbo went into the kitchen, then Gandalf came back out. (ruled in by PG&B) ✗
- (4) *ok* Bilbo ran to the Shire, then ran back. (ruled out by modified PG&B) ✗

What we really want is a way to say that the events are reversed, not just a single property. And there is simply no way to do that using properties. If we modify the definition and propose a macro-property $\mathcal{P} = \bigcup P \mathcal{P}(e)$, as reverse to $\mathcal{P}(e')$, something else goes wrong. We end up ruling out the core case of Source/Goal reversal in (4), where all other properties are identical e and e' .

SOLUTION I propose the abstract adverb COUNTER has the denotation (5). Source/Goal reversal and grammatical role reversal (both from P-G&B 2017) are reframed below to make them all properties of events (two-place function TRACE from Zwarts 2005; where 0 means

start point and 1 means end point). This works in combination with a condition forcing identity (6).

(5) RELATION BETWEEN EVENTS

- $\llbracket \text{COUNTER} \rrbracket = \lambda P_{st} \lambda e_s. \left[\exists e'_s \exists \mathcal{R}_C [\tau(e') \prec \tau(e) \wedge \mathcal{R}_C(e)(e')] \right] . P(e)$, where
 $\mathcal{R}_C = f \in D_{s,st} . \lambda e_s \lambda e'_s . f(e)(e') = 1$ iff either *Case 1* is true or *Case 2* is true
- Case 1: $\text{TRACE}(e)(0) = \alpha$, $\text{TRACE}(e)(1) = \beta$, and
 $\text{TRACE}(e')(0) = \beta$, $\text{TRACE}(e')(1) = \alpha$
- Case 2: $\text{Agent}(e) = \mathbf{x}$, $\text{Patient}/\text{Recipient}(e) = \mathbf{y}$, and
 $\text{Agent}(e') = \mathbf{y}$, $\text{Patient}/\text{Recipient}(e') = \mathbf{x}$

(6) IDENT-PRESUPPOSITION

Don't presuppose anything you can't get from the sentence itself.

ADVANTAGES *First:* Stating the meaning this way, with an explicit identity condition (6), means it does not have the problem of non-identity of some properties. The only way you can be counterdirectional is if you meet one of two narrow reversal conditions, and the everything in the sentence that was not reversed, should be exactly the same. In (3), even though it is true there that e and e' are counterdirectional according to *Case 1*, they get ruled out because their non-identical Agents violate (6). So, (3) is correctly ruled out by the identity condition above, IDENT-PRESUPPOSITION.

Second: We can now *swap* properties, instead of simply saying the earlier existential statement. The previous definition of the adverb would rule in, for example, a case where the Agent of e is the same as the Patient of e' (so far so good!) but the Patient of e is not the same as Agent of e' , but some other entity (bad!). This is quite important, because that bad result looks like this:

- (7) *Context:* Bilbo hugged Frodo, Gandalf hugged Aragorn. Then,
- a. #Frodo hugged Gandalf back. (ruled in by PG&B) X
 - b. #Aragorn hugged Bilbo back. (ruled in by PG&B) X

From this I conclude that when we take into account the problems with predicate-reversal and the advantages of event-reversal, event-reversal wins.

SYNTAX (TIME PERMITTING) I argue in the paper that the origin of IDENT-PRESUPPOSITION is syntactic constraints on what can/cannot go into the presupposition. IDENT-PRESUPPOSITION can be violated in Hindi-Urdu, but only if the counterdirectional *vaapas* is structurally unusually high. In (9), the verb is non-identical but the higher-than-(8) position of *vaapas* makes it ok. This is discussed more in the paper as an unnoticed “in return” reading.

- (8) #bil-ne gaen-ko dhakelaa, gaen-ne us-ko vaapas maaraa
 Bil-ERG Gan-DOM push.PFV Gan-ERG 3P-DOM COUNTER hit.PFV
 INTENDED: ‘Bil pushed Gan, Gan hit him back.’
- (9) bil-ne gaen-ko dhakelaa, gaen-ne vaapas us-ko do thappaR maare
 Bil-ERG Gan-DOM push.PFV Gan-ERG COUNTER 3P-DOM two slaps hit.PFV
 Lit. ‘Bil pushed Gan, *in return* Gan hit him two slaps.’

Uniqueness and familiarity in interpreting definite descriptions

Sadhwi Srinivas and Kyle Rawlins

Johns Hopkins University

Introduction. According to uniqueness-based theories of definite descriptions (e.g., Russell 1905, Evans 1977, Lobner 1985), such descriptions are felicitous only in contexts in which there is a unique referent satisfying the literal content of the description. On the other hand, familiarity-based theories claim that uniqueness is not crucial for referential success as long as the referent has been made salient in the preceding discourse context (e.g., Christopherson 1939, Kamp 1981, Heim 1982). In this study, we experimentally investigate the process by which hearers interpret definite descriptions in English in order to help moderate between these two types of theories of definite descriptions. Following this, we describe a Rational Speech Acts-style model (RSA; Frank & Goodman, 2012) implemented to fit our experimental data, and discuss in what ways participants adhere to and diverge from rationality.

Experiment. 40 M-Turk read descriptions of ten situations (*stories*) imagining themselves as one of the characters (the *hearer*). At the end of each story, they were asked to interpret a definite description uttered by the character they were interacting with (the *speaker*). We employed a 2 (speaker uniqueness) x 2 (hearer uniqueness) x 2 (linguistic salience) within-subjects design, where **speaker uniqueness** was true if the speaker knew there to be a unique referent satisfying the literal content of the description, and **hearer uniqueness** was true if the hearer knew there to be a unique referent satisfying the description. **Linguistic salience** was our proxy for familiarity which held of an object if it had been mentioned by the speaker or the hearer prior to the critical description to be interpreted.

The stories were designed to give hearers an unambiguous view of both the situation itself, as well as the speaker's relevant knowledge state. In every situation, two objects featured prominently. The status of these objects varied independently on speaker uniqueness, hearer uniqueness and salience. Fig (1) shows a trial in which speaker uniqueness was true (the auditor knew only the labeled jar to be *the jar of camphor*) but hearer uniqueness was not (to the participant imagining themselves as the lab worker, both the labeled and the unlabeled jar were jars of camphor). Additionally, the unlabeled jar was made linguistically salient by the auditor. At the end of the description, participants playing the role of the hearer of the description were asked to choose a referent corresponding to the description. They could pick the labeled jar, the unlabeled one, or refrain from picking either by selecting “*Don't know*” instead.

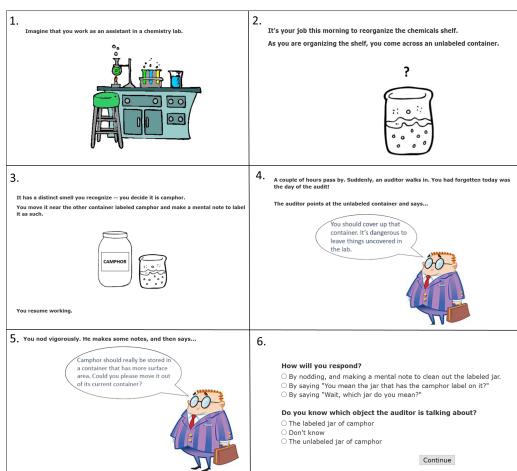


Figure 1: Example of an experimental trial

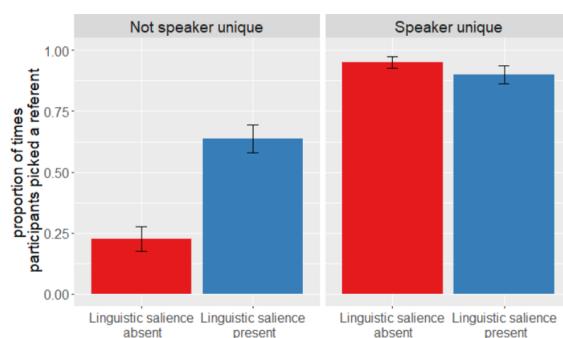


Figure 2: Uniqueness vs familiarity

Results. The main findings from our experiment are as below:

1. Participants primarily achieved referential success as long as speaker uniqueness held, regardless of the status of hearer uniqueness.

2. We observed ranked effects of uniqueness & familiarity as cues towards referential success. Speaker uniqueness led to referential success regardless of salience/familiarity (right panel in fig 2). In the absence of speaker uniqueness, salience led to ref success ~60% of the time (left panel in fig 2).
3. When the non-speaker-unique object (*distractor*) was salient (the unlabeled jar in the example trial), people tended to choose it more often (~20%) than they did a non-salient distractor (~8%).

Computational Model. The RSA explicitly models how a conversational agent accounts for their interlocutor's knowledge state via iterative Bayesian reasoning within a communicative task. Our experimental set-up is quite suited to be modeled within such a framework, especially given that participants reasoned mostly from the speaker's point of view. Our implementation of the RSA builds in three ways on the traditional implementations discussed in the literature:

1. Typically, the RSA model is employed in situations where the literal semantics are shared by both speaker and hearer. But assuming such a shared semantics wasn't appropriate in our scenarios which were explicitly designed to provide hearers with privileged knowledge that wasn't shared by the speakers. To handle this within the model, we assumed that hearers started with a speaker-centered semantics (as the experimental results indicated); however, noise could be introduced by the hearer's knowledge state. This noise measure was included as a trained parameter.
2. We operationalized salience by incorporating it into the prior. This implements the intuition that more salient objects have a higher chance of being talked about.
3. Finally, traditional implementations of the RSA do not give a straightforward way to refrain from choosing an alternative; however this action was allowed in our experiment ("Don't know"). To decide when to fail to choose a referent, we employed the idea that reference failure was proportional to the entropy of the rational listener's belief distribution. The entropy is translated to a probability value by a logistic function, and the probabilities of picking the referents modulo failure are then recomputed accordingly.

When we trained the model on our data, we found that the patterns of referential choice it predicted differed in certain meaningful ways from what was empirically observed. The model correctly captures that salience doesn't add too much on top of speaker-uniqueness when the latter holds (indicated in green in fig 3). But there are 2 main behavioral patterns it fails to capture. First, the model underestimates the extent to which hearers are susceptible to salient distractors, by failing to pick them even nearly as often as people do (indicated in red). Second, it overestimates the extent to which salience of a potential referent acts as a cue for referential success in the absence of speaker-uniqueness (indicated in orange).

Situation	RSA model			Actual data		
	p(A)	p(B)	p(failure)	p(A)	p(B)	p(failure)
A satisfies the literal content of the description; B does not; neither A nor B is salient	0.83	0	0.17	0.8625	0.0875	0.05
A satisfies the literal content of the description; B does not; A is salient	0.83	0	0.17	0.8625	0.0375	0.1
A and B both satisfy the literal content of the description; neither A nor B is salient	0.175	0.175	0.65	0.1	0.125	0.775
A and B both satisfy the literal content of the description; A is salient	0.71	0.03	0.26	0.55	0.0875	0.3625
A satisfies the literal content of the description; B does not; B is salient	0.63	0.05	0.32	0.6	0.23	0.175

Figure 3: Comparing model predictions to actual experimental data

Conclusion. The contribution of the current study is two-fold. First, we empirically evaluated the relative contributions of uniqueness and familiarity in the interpretation of definite descriptions. We found ranked effects of uniqueness and familiarity, where uniqueness acts as the stronger cue. Then, we implemented an RSA model to fit the experimental data and found that although participants behaved "rationally" in many ways, they were also more susceptible to salient distractors than the RSA predicts.

References. Russell 1905, *Mind*. Evans 77, *Canadian journal of phil.* Lobner 85, *J of S.* Kamp 81, *Formal methods in the study of lang.* Heim 82, UMass PhD. Frank & Goodman, 2012, *Science*.

Agent entailments within the semantics of roots

Josep Ausensi (josep.ausensi@upf.edu)

Universitat Pompeu Fabra

Introduction. Theories of event structure (Lexicalist: Rappaport Hovav & Levin 1998; Constructionist: Goldberg 1995; Neoconstructionist: Borer 2005, Ramchand 2008) assume that verb meanings are ‘bifurcated’ (Bifurcation Thesis (Embick 2009)) (BT) into *templates* (e.g. $\sqrt{\text{BECOME}}$) and *roots* (Pesetksy 1995) (e.g. $\sqrt{\text{BREAK}}$). In verbs like *widen* the entailment of change is introduced templatically, e.g. when the root $\sqrt{\text{WIDE}}$ is placed under $\sqrt{\text{BECOME}}$. Roots then never introduce templatic meanings, they only provide encyclopedic knowledge. Nonetheless, Koontz-Garboden and Beavers (2016) (KGB) argue that, contrary to the BT, roots like $\sqrt{\text{BREAK}}$ (in contrast to $\sqrt{\text{WIDE}}$) always introduce entailments of change (cf. #*The broken chair has never been broken* vs. *The wide table has never been widened*), and therefore KGB argue that such templatic meaning must come from the root itself, contra the BT (cf. *widen*).

Backdrop. (a) Kratzer (1996) (also Bowers 1993; Chomsky 1995; Collins 1997) argues that external arguments are not part of the verb’s argument structure as they are introduced by (silent) light verbs (e.g. little v). Kratzer (following Marantz 1984) points out that only objects are ‘true’ arguments since verbs only impose selectional restrictions on objects, e.g. for *kill* to have the interpretation of ‘spend time doing x ’ (e.g. *kill an afternoon reading books*) it selects an object that must denote time intervals, i.e. idiomatic meanings are only triggered by internal arguments. Kratzer notes that subjects, on the other hand, are rather special since verbs do not impose selectional restrictions on them, and therefore the type of subject rarely alters the meaning of the verb (but see Nunberg, Sag and Wasow 1994). While this analysis has been generally adopted (Marantz 1997; Pylkkänen 2002; Harley 2013, *i.a.*), it has been challenged and remains controversial (Kiparsky 1997; Krifka 1999; Horvath and Siloni 2003; Wechsler 2005). **(b)** More recently, Folli and Harley (2005) (FH) argue that a new typology of little v is needed in order to account for different subject types, i.e. consumption verbs (e.g. *eat*) only allow animate subjects, yet nonanimated subjects are permitted in resultative-like constructions (cf. *The sea ate *(away) the beach*). FH posit then different ‘flavors’ of v to account for such variation, i.e. v_{CAUSE} and v_{DO} . Subjects, thus, can either be assigned an Agent or a Cause role depending on the ‘flavor’ of v : v_{DO} requires that the subject qualify as an Agent (it must be intentional) whereas v_{CAUSE} only requires that the subject be capable of initiating a change. As FH point out, these flavors can be lexically encoded, (e.g. *murder*) (see also Van Valin and Wilkins 1996). Nonetheless, the fact that verbs like *murder* impose selectional restrictions on their subjects (cf. #*John accidentally*/#*The weapon murdered Tom*) is unpredicted by Kratzer’s analysis insofar as external arguments are never introduced by verbs. FH suggest that it is the ‘vocabulary entry’ for verbs like *murder* what requires external arguments to appear under v_{DO} .

Proposal. Following KGB, I argue that **(a)** some roots come with the templatic meanings introduced by v_{DO} (i.e. intentionality/agent entailments) contra the BT. I then argue **(b)** against Kratzer (1996, *i.a.*) that verbs do not impose selectional restrictions on external arguments. In this respect, I claim that the $\sqrt{\text{MURDER}}$ root class ($\sqrt{\text{MURDER}}$, $\sqrt{\text{ASSASSINATE}}$, $\sqrt{\text{SLAY}}$, $\sqrt{\text{BUTCHER}}$, $\sqrt{\text{MASSACRE}}$) have as part of their entailments the templatic meanings introduced by v_{DO} , contra the BT, with grammatical consequences.

Analysis. Following Beavers and Koontz-Garboden (2018: 159), I propose the denotation in (1) for $\sqrt{\text{MURDER}}$ roots: (1) differs from the one that Beavers and Koontz-Garboden propose for roots like $\sqrt{\text{GUILLOTINE}}$ in that the cause that brings about the change in $\sqrt{\text{MURDER}}$ only needs to be of an intentional-type action (i.e. performing an action intentionally to bring about the result state of *dead*). (1) predicates the result *dead* of a unique argument but it specifies that such result state must have a specific cause. (1) does not apply to roots like $\sqrt{\text{KILL}}$.

$$(1) \lambda x \lambda s [\text{dead}'(x, s) \wedge \exists e' \exists v [\text{cause}'(v, e') \wedge \text{become}'(e', s) \wedge \forall v' [\text{cause}'(v', e') \rightarrow \text{intentional}'(v')]]]$$

Evidence. (i) $\sqrt{\text{MURDER}}$ roots cannot be ‘disentangled’ from the meanings introduced by v_{BECOME} and v_{DO} , i.e. from the entailments that there was a change (KGB) (2a) and was brought

about by a cause that is intentional and “can *do* things” (FH: 110) (2b)-(3a), suggesting that such entailments must be part of the meaning of the root itself. This contrasts with roots like $\sqrt{\text{KILL}}$: v_{DO} meanings can be ‘disentangled’ (4a) since $\sqrt{\text{KILL}}$ is underspecified for the type of cause and therefore allows different types of causes (4b) in contrast to $\sqrt{\text{MURDER}}$ roots (3b).

- (2) a. #The murdered/assassinated president hasn’t been murdered/assassinated.
- b. #The murdered/assassinated president hasn’t been killed intentionally.
- (3) a. #Bob slew/butchered the dragon but didn’t do it on purpose/but wasn’t his intention.
- b. #The sword/#The arrow/#The poison slew/butchered the dragon.
- (4) a. The sniper killed the president but didn’t kill him intentionally (=it was an accident).
- b. The poison/The flood/The arrow/The sword/The guillotine killed the president.

(ii) Crucial evidence comes from *again* modification (McCawley 1968; Dowty 1979; von Stechow 1996): with $\sqrt{\text{MURDER}}$ roots, *again* always scopes over both meanings, i.e. the specific cause (v_{DO}) and the result (v_{BECOME}) (6). This contrasts with resultative constructions where *again* can scope just over the result (5) (see Beck and Snyder 2001; Beck and Johnson 2004).

- (5) Mary made a sheet of metal that is flat, but it later accidentally became bent. Fortunately, John hammered the metal flat again. (Beavers & Koontz-Garboden 2012: 357)

The reading in (5) is restitutive since the metal does not need to have been hammered in a previous stage or even flattened (Beavers & Koontz-Garboden 2012: 357) since *again* in (5) scopes just over the result ([[John DO < $\sqrt{\text{HAMMER}}$ >] CAUSE [the metal BECOME< $\sqrt{\text{FLAT}}$ >]]). In contrast, *again* modification with $\sqrt{\text{MURDER}}$ roots always scopes over both meanings as they are ‘packaged’ together in a single root (Beavers & Koontz-Garboden 2012): with $\sqrt{\text{MURDER}}$ roots *again* cannot scope over just the result as with resultative constructions with two different roots (5). Further, *re-* modification (7) (Marantz 2007; Beavers and Koontz-Garboden 2012) yields the same readings as *again* modification. This shows that the denotation in (1) for such roots appears to be correct.

- (6) He murdered/slew/massacred/butchered the monsters again.

MEANS	He caused the monsters to become dead by <u>intentionally</u> killing them again.
CANNOT MEAN	He caused the monsters to become dead again by intentionally killing them but the last time they were killed <u>by accident/unintentionally</u> .

- (7) He remurdered/reslew/remassacred/rebutchered the monsters.

MEANS	He caused the monsters to become dead by <u>intentionally</u> killing them again.
CANNOT MEAN	He caused the monsters to become dead again by intentionally killing them but the last time they were killed <u>it was not due to intentional-type actions</u> .

(iii) The present analysis also accounts for the cross-linguistic fact that verbs derived from $\sqrt{\text{MURDER}}$ roots never alternate between causative/inchoative uses (e.g. **Kennedy assassinated*). This is predicted since $\sqrt{\text{MURDER}}$ roots introduce (templatic) meanings that require a specific cause and, as it has been noted, verbs placing such restrictions on subjects never alternate (Koontz-Garboden 2009; Schäfer 2009; Rappaport Hovav 2014; Alexiadou, Anagnostopoulou & Schäfer 2015, *i.a.*). In contrast, verbs derived from roots lacking v_{DO} templatic meanings can alternate ($\sqrt{\text{KILL}}$), e.g. *kill* (in Spanish (Vivanco Gefaell 2016), Hebrew (Rappaport Hovav 2014) or Greek (Schäfer 2008)), which is also predicted under the present account. Thus, insofar as $\sqrt{\text{MURDER}}$ roots come with the templatic meanings introduced by v_{DO}, such roots are always lexicalized as causative verbs since, as Beavers and Koontz-Garboden (2018: 61) note, roots may be “directly lexicalized as causative verbs if specified for the cause.”

Conclusion. I argue that $\sqrt{\text{MURDER}}$ roots come with the templatic meanings introduced by v_{DO}, which result from the meaning of the root itself, contra the BT, with grammatical consequences. This accounts for the fact that verbs derived from $\sqrt{\text{MURDER}}$ roots place selectional restrictions on their subjects (contra Kratzer 1996, *i.a.*), never have inchoative uses and cannot be ‘disentangled’ from the templatic meanings introduced by v_{DO}.

Direct vs. indirect causation: A new approach to an old problem

Rebekah Baglini (Aarhus University) and Elitzur Bar-Asher Siegel (Hebrew University of Jerusalem)

Introduction English causative constructions come in simplex and complex variants, exemplified in English by *Sam killed Lee* and *Sam caused Lee to die*. While use of the former, the lexical causative, entails the truth of the latter, an entailment in the other direction does not hold. This difference is often ascribed to the lexical causative having an additional prerequisite of **direct causation**, such that the causative relation holds between a contiguous cause and effect, and no third event is allowed to intervene (Fodor (1970); Katz (1970); Wolff (2003), *inter alia*). However, capturing direct causation formally is a challenge (Dowty, 1979), especially in light of paradoxical empirical patterns. For example, the lexical causative in (1) involves causal chains where cause and effect are not contiguous.

- (1) Opening bus lanes to motorcycles will redder the streets of London with cyclists' blood.
[*opening bus lanes* > *accidents increase* > *some cyclists die*]

This paper aims to formally capture the causative components of lexical causative verbs by applying insights from causal models for counterfactual reasoning. This analysis will be able to explain the intuitions that lead to the direct causation hypothesis, as well as to account for counterexamples like (1).

Previous work: Bar-Asher Siegal and Boneh (2018) propose that lexical causatives presuppose that the cause was sufficient for the effect, while cause asserts that it was only necessary. To address the problem of directness, the current paper enriches the notion of sufficiency through the use of Structural Equation Models (SEMs) (Pearl, 2000; Schulz, 2011) for causation. Such models have been applied previously in analyses of the verbs *cause* and *make* (Lauer and N., 2017) among other periphrastic constructions (Baglini and Francez, 2016).

Formal machinery: The truth conditions of causal statements rely on a structural causal model. Accordingly, semantic accounts provide a definition for when a causal statement is true, given such models. Dependencies between states of affairs are represented in a SEM by a set of propositions (variables) and truth values. To illustrate, consider the variables involved in an opening of an electric door:

A. Electricity: =1 if running; else =0	D. Door opens: =1 if opens; else =0
B. Button: =1 if pressed; else =0	E. Button =1 \models Circuit =1
C. Circuit: =1 if closed; else =0	F. Circuit =1 & Electricity =1 \models Door opens =1

The fact that some variables depend on others for their value is represented by structural equations, in E. and F. We use the term **condition** for a variable that necessarily determines the value of another. We adopt the “dynamics” framework of Schulz (2011) to identify both the set of immediate causal ancestors for a proposition, as well as the nature of the direct dependencies. A world is defined as a model with all variables valued. In this context, a **causal basis** for a possible world w and a dynamics D is defined as the minimal set of facts in w which determines everything else in w . Thus, in the above example, the basis consists of Button and Electricity: the values of all other variables are determined based on the values of these two conditions. In our analysis, we recast sufficiency in terms of a causal basis: the minimal set (possibly singleton) of conditions that are sufficient for the effect. Thus, unlike previous analyses, sufficiency here is not restricted to a single condition, but is rather a characteristic of a set of conditions (Cf. Mackie (1965)), defined here as the basis.

Analysis: Linguistic expressions of causation commonly presuppose models with the above characteristics, adding construction-specific entailments and pragmatic inferences. For example, periphrastic *cause* selects as its subject any condition on which the value of the effect causally depends—i.e. any necessary condition (2) (Lauer and N., 2017).

(2) $\{\text{Electricity}/\text{John} \ (\text{pressing of the button})/\text{the closed circuit}\}$ causes the door to open.

In contrast, lexical causatives are sensitive to event-related changes in value of conditions in the basis, given that events have time in which the value of associated variables in the model can be changed ($0 > 1$ or $1 > 0$). It is possible, then, to identify the event which “completes” the basis, such that following this event (but not before) the values of the set of conditions in the basis entail that the effect occurs (occurrence is also defined as change of values). A lexical causative asserts that its subject is a participant in this **completion event**. Considering the context of an automatic door, this event *usually* corresponds to the Button condition (3-a). Thus, (3-b), which substitutes Electricity, is ruled out.

- (3) a. $\{\text{Pushing the button}/\text{John's pushing the button}/\text{the button}/\text{John}\}$ opened the door.
 b. Electricity opened the door.

The notion of a completion event is needed to explain why judgments favor (3-b) over (3-a) when the door scenario is changed such that the button is depressed beginning at time $t-1$, but a power outage prevents the door from opening until electricity is restored at time t . This alternative context involves two non-simultaneous event-related conditions in the basis (Button and Electricity). When the temporal order of events is retrievable contextually, the lexical causative must express the completion event—i.e. Electricity (3-b). The causative component of a lexical causative verb is represented formally in (4). F is a function which takes the value of all the conditions in the basis at a given time t . It returns 1 if the values of the basis set entail the occurrence of the effect and 0 if they do not. AT is a Davidsonian predicate mapping events to their temporal envelope.

(4) $\exists e \exists t : P(e) \ \& \ AT(t, e) \ \& \ P \in \text{BASIS} \ \& \ F(t, \text{BASIS}) = 1 \ \& \ \forall t' \prec t \forall e' : AT(t', e') \rightarrow [\neg P(e')]$

The formula in (4) amounts to a description of a **completion event**: a condition P is part of the set of conditions that constitutes the basis. At the time t of the event affecting the value of P (i.e. prior to it, for all events, $\neg P$), the model determines that the occurrence of the effect must take place. Since prior to t the value returned by F had to be 0 (since P is part of the basis set), the event at time t is the completion event.

Back to direct causation: Contrary to the directness hypothesis, lexical causatives do not require contiguity between cause and effect at all. The intuition of direct causation arises (epiphenomenally) from contrasting lexical causatives with periphrastic *cause*: the stronger selection pattern of the former may exclude more temporally distant conditions, while the latter admits any necessary condition. This gives the illusion of a stronger contiguity requirement for lexical causatives. Our analysis explains the acceptability of sentences like (1) which involve events intervening between the subject-affiliated event and the effect, when each intervening condition is understood to be fully determined from the basis. In example (1), the increase in the number of accidents is not in the basis as it depends on (i.e. is a necessary result of) the opening of the bus lanes. Thus, this paper takes the longstanding problem of direct causation inferences and uses models for counterfactual reasoning to explain why lexical causatives selectively express directness.