Agentive meaning

In his foundational 1957 paper, Grice introduced a distinction between natural and non-natural meaning, using a series of linguistic tests to argue that non-natural meaning is all and only meaning with an agent. This observation underpins the entire Gricean project, namely that we should consider speaker meaning as primary to utterance meaning. The goal of this paper is to argue that Grice's (1957) tests only indirectly and imprecisely test for (non-)natural meaning; they instead reflect tests for agentivity and eventivity. And while natural meaning is necessarily non-agentive, non-natural meaning can be either agentive or non-agentive. While this doesn't undermine the Gricean project, it does weaken his empirical claim, and makes the right predictions for languages that appear to lexically differentiate between natural and non-natural meaning.

I begin the paper by reviewing clear counterexamples to this biconditional (some of which Grice noted himself): examples of intuitively non-natural meaning that lacks an agent (entity). In §2–2.2 I examine Grice's linguistic tests through the lens of modern-day lexical and event semantics, and argue that, instead of tracking (non-)naturalness, they track two independent but related linguistic properties: agentivity (whether or not an agent is encoded in the syntactic argument structure) and eventivity (whether the eventuality argument associated with the meaning relation is an event or a state). In §3 I explain how this new perspective provides evidence for a weakening of Grice's biconditional: while natural meaning is necessarily non-agentive and stative, non-natural meaning can come in any variety. Finally, I take up Grice's question of whether the verb *to mean* is polysemous, ambiguous, or neither, by reflecting on the nature of near-synonymous verbs in English (§4.1) and the behavior of synonyms and near-synonyms cross-linguistically (§4.2).

But first, I'll present Grice's original distinction, and the tests he used to justify it.

1 Grice's (non-)natural meaning distinction

In discussing the meaning relation, I will use the following terms and variables:

(1) a. x is the **signaler** (or speaker, for spoken languages)
b. σ is the **signal** (a linguistic string, sign, etc.)
c. μ is the **meaning** associated with a given signal

The discussion in Grice (1957), and a vast subsequent literature, takes for granted that we have intuitions about which meaning relations are (non-)natural. The goal of this paper is not to underpin or define (non-)natural meaning, but to instead argue that the distinction is not linguistically marked, as Grice argued, but is instead epiphenomenal, more closely correlated with distinctions of agentivity and eventivity. And, for better or worse, a paper whose goal it is to argue that Grice was mistaken in thinking (non-)natural meaning is linguistically marked must take for granted Grice's assumptions about what qualifies as natural and non-natural meaning.

The sentences in (2) and (3) are exemplifications of the intuitive (non-)natural meaning distinction, lifted or amended from Grice's original discussion.

(2) **natural meaning**
a. Those spots mean measles.
b. The recent budget means that we will have a hard year.

(3) **non-natural meaning**

a. That bell ring means the bus is full.
b. 'Knucklehead' means Peter.

According to Grice, the sentences in (2) qualify as natural meaning because they report a stable, non-contingent relationship between the signal and the signified. In contrast, the sentences in (3) qualify as non-natural meaning because they report a contingent relationship between the signal and the signified. Despite Grice's presumptions that the sentences in (2) and those in (3) exemplify two different meaning relations, there is nothing superficially linguistically different between them (i.e., both are of the form 'σ *mean*(*s*) μ'). In other words, Grice's expectation is that, independently of their knowledge of their language, speakers are bringing to the table a distinction between (2) and (3), and he goes on to highlight what he takes to be linguistic reflexes of this distinction.

In his support of the (non-)natural meaning distinction, Grice presents five linguistic diagnostics. The labels are mine; Grice (1957) doesn't explicitly discuss e.g. passives or (linguistic) agents, just the examples he believes qualify as linguistic diagnostics for (non-)natural meaning.

Passivizability All and only sentences of the form 'σ means μ' that convey non-natural meaning (e.g. (3), but not (2)) entail a sentence of the form 'What was meant by x is σ'.

(4) a. ??What was meant by those spots was that he had measles.
b. What was meant by that ring was that the bus is full.

This is one test for which the distinction between x subjects (where x ranges over signalers, or speakers) and σ subjects (where σ ranges over signals, or linguistic strings) is important. The passivizability test is crucially a test of sentences with the form 'σ means μ', which is the form of all four sentences in (2)-(3).

Agentivity All and only sentences of the form 'σ means μ' that convey non-natural meaning (e.g. (3) but not (2)) entail a sentence of the form 'Somebody or other meant μ by σ'.

(5) a. ??Somebody meant by those spots that he had measles.
b. Somebody meant by that ring that the bus is full.

This is clearly related to the passivizability test above: intuitively, we can interpret the verbal passive μ *was meant* as entailing a syntactic agent, just as the sentence *Mistakes were made* entails that there was someone who made the mistakes. I discuss this in more detail in §2.

Actuality All and only sentences of the form 'σ means μ' that convey natural meaning (e.g. (2), but not (3)) commit the speaker to μ being actual.

(6) a. Those spots mean measles, #but he hasn't got measles.
b. That bell ring meant the bus is full, but it isn't in fact full.

There are a few reasons I avoid the terms 'factivity' and 'veridicality' here (cf. Neale, 2022). First, the term 'factivity' is most closely associated with attitude predicates, and I want to separate for now any kinship between *means* and e.g. *knows*.[1] Second, while Grice's description of this diagnostic dovetails with how we typically characterize the phenomena of factivity and veridicality – as entailment to a proposition, either writ large or relative to the speaker's beliefs – this isn't a necessary property of meaning constructions (either natural or non-). In other words, Grice characterizes this test as follows: "That is to say, in cases like [(2)], '[x] meant that p' and '[x] means that p' entail p" (p377). But there are plenty of cases where the signified is not a proposition, as in (2-a). Preserving this factivity or veridicality notion in the existence of signifiers like the noun *measles* would require a sophisticated theory of ellipsis or some counterpart that, for any context of utterance, transforms the eventuality-signifying argument *measles* (see §2.2.2) into a proposition. I do not have such a theory, and so am opting for a less-loaded term here.

[1] I will return to discuss this potential kinship in §4.1.

Direct Quotability All and only sentences of the form 'σ means μ' that convey non-natural meaning (e.g. (3), but not (2)) can be paraphrased using quotes for the μ argument.

(7) a. *The recent budget means 'We will have a hard year'.
b. Those three rings on the bell mean 'The bus is full'.

This test is a tricky one; it is hard to generalize the classic notion of a linguistic intuition to orthographic conventions. If we consider quotation marks to reliably mark quotation intonation (Sturman, 2021), we can interpret Grice's Direct Quotability test as a test of our intuitions about when we can insert an intonational break immediately before the μ string, and whether μ can be pronounced with a pitch range reset.

This is a markedly more subtle linguistic judgment than the others, and I empathize with those who are tempted to throw their hands up and abandon this test. However, thanks to phonetic work like Sturman (2021), I do believe in the prosodical reality of quotation, and I will proceed as though this is a reliable test, based on my own (admittedly weak) intuitions about Direct Quotability.

Factiness In all and only sentences conveying natural meaning can the signal σ be truth-preservingly paraphrased as a fact (e.g. in sentences like those in (2), but not like those in (3)).

(8) a. The fact that he had those spots meant that he had measles.
b. #The fact that the bell has been rung three times means that the bus is full.

Following the work in Grice (1957), there was an important if short-lived focus on the semantics of factiness, i.e. the distribution of *the fact that* clausal marking (Ginzburg and Sag, 2001). These judgments are also subtle; I and an anonymous reviewer, for instance, find (8-b) to be completely acceptable. But for those who agree with Grice's characterization of the data, I will argue that it dovetails in a particular way with the Direct Quotativity test.

Table 1 summarizes Grice's five linguistic diagnostics for the (non-)natural meaning distinction.

	ex.	*description*	**natural?**	**non-?**
Passivizability (a)	(4)	Can 'σ means μ' be passivized?	✗	✓
Agentivity (b)	(5)	Does 'σ means μ' entail 'Somebody means σ'?	✗	✓
Actuality (c)	(6)	Is the speaker committed to μ being actual?	✓	✗
Direct Quotability (d)	(7)	Can the signified σ argument occur in quotation marks?	✗	✓
Factiness (e)	(8)	Can the signifier s be paraphrased with 'the fact that'?	✓	✗

Table 1: Grice's (1957) linguistic diagnostics for the (non-)natural meaning distinction

My claim here is that there is a tight but imperfect correlation between the (non-)natural meaning distinction and two independent linguistic properties: agentivity (whether or not an agent is encoded in the syntactic argument structure) and eventivity (whether the eventuality argument associated with the meaning relation is an event or a state). The latter will also get tied up a bit in tests for whether a signal σ denotes an actual event.

I argue that, while natural meaning is necessarily non-agentive and stative, non-natural meaning can come in any variety, so there are meaning constructions for which Grice's distinction and mine fall apart. But I first present the data discussed above from a more modern and informed linguistic lens, so we can better control for what we are diagnosing, and thereby have a firmer grasp on the significance of our conclusions.

2 The semantics of meaning constructions

2.1 The morphosyntax of meaning constructions

Verbs divide into classes with respect to the argument alternations they allow (Levin, 1993). Two examples of argument-alternation classes are as follows:

(9) a. The sun melted the ice.
b. The ice melted.

(10) a. The sunlight changed the color.
b. The color was changed by the sunlight.

The ergative verb *to melt* exemplifies the causative alternation of the class of unaccusative verbs; it can assign a theta-role to an agent subject and a theme object (as in (9-a)). But the verb *to melt* has another argument realization pattern, e.g. (9-b), in which the semantic theme can occur as the verb's surface subject. This is not true of unergative verbs like *met* (cf. *Magda met Matthew* and **Magda met*).

To change in (10) exemplifies a different argument-alternation class of verbs whose causer subject (*the sunlight* in (10-a)) and theme object (*the color* in (10-a)) can effectively switch places, linearly. Unlike ergative verbs, the derived form of these unaccusatives can specify its causer, as in (10-b), but it can no longer be a syntactic argument of the verb; it must be introduced in a *by*-phrase (*by the sunlight*).

I believe it is a novel observation that *to mean* exemplifies (in a way) an even quirkier class of verbs, of which *to cover* and *to amuse* are also members (Rappaport Hovav, 2024).[2] Each of these verbs exhibit three argument realization options, which are exemplified below and schematized (for *to mean*) in (14).

(11) a. The blanket covered the screen.
b. Una covered the screen.
c. Una covered the screen with the blanket.

(12) a. Her antics amused Adam.
b. Karen amused Adam.
c. Karen amused Adam with her antics.

(13) a. 'Knucklehead' meant Peter.
b. Ernie meant Peter.
c. By 'Knucklehead', Ernie meant Peter.

(14) a. σ meant μ. *unaccusative*
b. x meant μ. *agentive*
c. By σ, x meant μ. *triadic*

All of the above argument alternations involve some combination of three arguments: (i) an agent x that performs an action; (ii) a theme σ that is used to carry out the action; and (iii) a μ that receives the action (but does not change state), varying from a location (in (11)) to an experiencer in (12).

The labels for the alternations in (14) are important. Unaccusative argument alternations (14-a) lack agents, even in some sort of deep-structure construal of the verb, and don't passivize, as illustrated in (15).

(15) *Peter was meant by 'Knucklehead'.

In contrast, the agentive form in (14-b) has an agent, but lacks a signal argument σ. Finally, the triadic form in (14-c) is the only one to include all three semantic arguments.

There is plenty of discussion about the relationship between the lexical semantics of a verb and its argument structure, and in particular which arguments are introduced where, and which are optional (see especially Levin and Rappaport-Hovav, 1995). I know of no such work that discusses *to mean* in particular, although in §4.1 I mention work on related phenomena in related classes of verbs. But this way of construing the argument alternations of *to mean* is compatible with the unaccusative version in (14-a) being its underlying form, because

[2]Thanks to Ethan Poole (p.c.) and Beth Levin (p.c.) for help here. The cover verbs actually come from two of Levin's (1993) classes: *fill* verbs – a subclass of verbs of putting – and verbs of contiguous location, a subclass of verbs of location. There remain subtle morpho-syntactic differences between *to cover*/*to amuse* on the one hand and *to mean* on the other that I am not able to probe deeply here (e.g. the difference in status of the argument introduced by a preposition, and which preposition is required to introduce it). *To mean* also shares some similarities in argument structure with the verb *to say* Grimshaw (2015), which I will return to in §2.2.1.

there is an asymmetric entailment relation between the agentive form and the unaccusative form.[3]

In sum, the verb *to mean* is relatively unusual in terms of its argument structure: it is a member of a class of only a handful of other verbs in English that have not just one but two optional semantic arguments, depending on the syntax of the construction (Rappaport Hovav, 2017; Wilson, 2020). And while it's helpful to think of *to mean* as relating an agent x with its signal σ and its signified meaning μ, there seems to be a natural argument that it is a verb with two basic arguments, σ and μ, and an optional agent (or signaler) x.

2.2 The (aspectual) semantics of meaning constructions

2.2.1 A stative/eventive alternation

The goal of this section is to review the stative/eventive distinction, and show that *to mean* is a lot like verbs of saying in its ability to have either a stative or eventive interpretation, depending on what sorts of things it's relating (and how).

Verbs can be stative or eventive, or either (Vendler, 1957; Dowty, 1979; Moens and Steedman, 1988). The distinction pertains to whether the verb characterizes an eventuality that is a state of being (a property of the subject), or an event. Unambiguously stative verbs are generally incompatible with the progressive and imperative mood in English (cf. **Elmar is knowing Icelandic* and **Know Icelandic!*), while eventive words (which include activities, achievements, and accomplishments), can do either or both.[4]

There has been some recent work on the behavior of the verb *to say* that will be instructive for the present study of *to mean*. Grimshaw (2015); Anand et al. (2017) and Major (2024) observe that *to say* usually takes a speaker as its subject, as in *Jane said that syntax is fun*, and when it does, these constructions are eventive: i) they can occur in the progressive aspect (16-a); ii) they can take a goal argument (i.e., a *to*-phrase) (16-b); and iii) they cannot occur in the simple present without receiving a habitual interpretation (16-c), amounting to a plurality of events distributed regularly over time.

(16) a. Jane was saying that syntax is fun (when Travis interrupted).
b. Jane said to Travis that syntax is fun.
c. Jane says that syntax is fun.

But they also observe that verbs of saying have a second 'variant' (in Grimshaw's terms), where they take a non-agentive subject, as in *The sign said that the park was closed*.[5] This variant of the verb, they argue, is stative as opposed to eventive. As a consequence, it does not hold the three properties listed above: the non-agentive variant i) cannot occur in the progressive (17-a); ii) cannot take a goal argument (17-b); and iii) they can occur in the simple present without receiving a habitual interpretation (17-c).

(17) a. *The sign was saying that the park was closed (when Travis vandalized it).
b. *The sign said to Travis that the park was closed.
c. The sign says that the park is closed.

Another test for an eventive interpretation (as opposed to a stative one) is compatibility with event modifiers like *yesterday*. When event modifiers modify an eventive verb, they specify the timeframe of the event. In contrast, when event modifiers modify a stative verb, they trigger a 'lifetime effect': an implicature that the situation no longer holds. This contrast is illustrated in (18).

(18) a. Jane said yesterday that syntax is fun. *eventive*
b. The sign said yesterday that the park was closed. *stative*

Of particular interest to us is the fact that *to mean* is just like verbs of saying not in terms of its argument

[3] Thanks to a reviewer for specifying this perspective.

[4] A reviewer points out that cover verbs, discussed above, are an unusual exception to the claim that statives can't occur in the progressive: we can say things like *Snow is covering the bridge*. This seems to be true of *to mean* in certain circumstances, as discussed below.

[5] Grimshaw also calls these 'say-by-means' constructions.

structure but in terms of its ability to take an agentive or non-agentive subject (and in the fact that this alternation conditions eventivity/stativity). This is illustrated in (19) and (20), which are intended to parallel (16) and (17), respectively.[6]

(19) a. ?Ernie was meaning Peter (when he said 'Knucklehead').
b. Ernie means Peter when he says 'Knucklehead'.

(20) a. *Smoke was meaning fire (when the firefighters arrived).
b. Smoke means fire.

That said, *to mean* differs from *to say* in that the latter is more naturally eventive. Specifically, it seems as though *to say* has a default eventive reading, but can have a stative reading in certain contexts; while *to mean* has a default stative reading, but can have an eventive reading in certain contexts. (This spectrum of eventivity in verbs of signification is discussed in more depth in §4.1.) Evidence for this comes from the interpretation of the event modifier *yesterday* in the unaccusative and agentive forms:

(21) a. (By 'Knucklehead',) Ernie meant Peter yesterday. *eventive*
b. Smoke meant fire yesterday. *stative*

This pair is intended to parallel (18). In the agentive form in (21-a), *yesterday* is interpreted as specifying the event of meaning. In contrast, in the unaccusative form in (21-b), *yesterday* produces a lifetime effect (an epistemically implausible implicature that smoke no longer means fire).

In summary, *to mean* appears to be a verb that is systematically ambiguous between an eventive and stative interpretation, just like verbs of saying (Grimshaw, 2015; Major, 2024). Just like for the verb *to say*, whether or not the subject of a *to mean* sentence is agentive plays a role in conditioning the eventive and stative readings. *To say* is generally eventive, unless it has a non-agentive subject, in which case it is much more naturally stative. *To mean* is generally stative, unless it has an agentive subject, in which case it can be eventive.

2.2.2 Event-denoting arguments can be interpreted independently from the verb

In this section, I switch from discussing the aspectual properties of the verb *to mean* to discussing the properties of its signal argument σ. I observe that, in meaning relations with an agent, the signal is typically a linguistic string, a sign, or a written word. But in meaning relations without an agent, the signal can be an eventuality (i.e., a state or an event).

There are lots of ways nouns or DPs can denote an eventuality; in English, a typical way is an action nominal or gerundive nominal (Comrie, 1976; Baker and Vinokurova, 2009).

(22) a. I found my wallet. It took all afternoon. *verbal event anaphora*
b. (The) finding (of) my wallet was an inconvenience. It took all afternoon. *nominal event anaphora*

As demonstrated in (22-a), the pronoun *it* can be anaphoric to an eventuality, which we typically characterize as being introduced by main verbs (here, *found*). But certain types of subjects can denote eventualities, as illustrated in (22-b). It's important to note that, consistent with the Davidsonian tradition, these subjects are introducing eventualities in the semantics that are distinct from eventualities associated with the main verb. And they can be interpreted differently, too; in (22-b), the eventuality introduced by the subject is interpreted episodically, as a single telic event (the finding of the wallet). But the eventuality introduced by the main verb, *was*, is a state; it's predicating the property *inconvenience* of the wallet-finding event. So, one sentence, two eventualities, with distinct aspectual properties.

We see this same property in meaning constructions. In particular, in meaning constructions without agents, the subject signal can denote an eventuality instead of an individual. And, in such cases, the eventuality can be interpreted eventively even while the main verb means is interpreted statively. This is illustrated in (23), where the gerund subject denotes an event of face-scratching.

[6]The question mark on (19-a) signifies a disagreement in acceptability between the author and a reviewer; I believe, however, that the contrast in acceptability between (19-a) and (20-a) is consistent across speakers.

(23) a. Sam scratching his face means that he has a bad hand.
b. Sam scratching his face meant that he has a bad hand.

Notice that this event subject can be interpreted eventively even while the verb has a stative meaning; an accurate paraphrase of (23) is 'The event of Sam scratching his face has the (stative) property of meaning that he has a bad hand.' In contrast, (23-b) has the same event-denoting subject as does (25a), but because of its past-tense marking, the main verb in (23-b) also has an eventive interpretation. (23-b) is paraphrasable as 'The event of Sam scratching his face, at some point in the past, had the property of meaning that he has a bad hand.' In sum, there is a lot of complex aspectual semantics lurking below the surface in meaning constructions.

3 Weakening the biconditional

Grice argued that natural and non-natural meaning come completely apart in the diagnostics presented in §1 and concluded, in particular, that all and only non-natural meaning is agentive. In this section, I present some counterexamples to this claim, and argue for a weakening of this biconditional, in which natural meaning is necessarily agentless and non-eventive, but non-natural meaning can come in any variety.

3.1 Non-agentive non-natural meaning

The goal of this section is to show that there is an indirect, but not one-to-one correlation between non-natural meaning, on the one hand, and agentivity and eventivity on the other; and to argue that, where the two fall apart, Grice's tests track agentivity and eventivity. In §4, I revisit Grice's question of the nature of the lexical semantics of *to mean*, and whether it makes sense to consider natural and non-natural meaning as two related interpretations of the same verb, or as encoded in two separate but homonymous verbs.

In contrast to Grice's (1957) claims, there are instances of non-natural meaning being non-agentive, and there are instances of non-natural meaning being stative. Grice himself discusses a counterexample:

> "I feel inclined to say that (as regards traffic lights) the change to red meant$_{NN}$ that the traffic was to stop; but it would be very unnatural to say, 'Somebody (e.g., the Corporation) meant$_{NN}$ by the red-light change that the traffic was to stop."' (p385).

Grice continues, "Nevertheless, there seems to be *some* sort of reference to somebody's intentions." There are two separate and separable claims here: what a given use of the verb *to mean* encodes and entails; and what our real-world knowledge tells us must be the case. In this discussion of traffic lights, the sentence in (24-a) is an unaccusative form of the verb; and in the absence of an agent x, the sentence is non-agentive (as per Grice's observation about the Corporation) as well as stative (it's incompatible with the progressive in (24-b), and the use of an event modifier in (24-c) triggers a lifetime effect).

(24) a. A red light means stop.
b. *A red light is meaning stop.
c. Yesterday, a red light meant stop.

It is, however, interesting to note that, when (non-)naturalness and (non-)agentivity fall apart, Grice's tests track (non-)agentivity instead of (non-)naturalness. Imagine two distinct scenarios of Sam playing poker. In Scenario A, Sam has a tell: he subconsciously scratches his cheek whenever he has a bad hand. (Scenario A is intended to be a case of natural meaning.) In Scenario B, Sam has a co-conspirator who is helping him cheat. He scratches his cheek as a signal to his co-conspirator that he needs help winning the hand. (Scenario B is intended to be a case of non-natural meaning.)

There are certain syntactic configurations that are compatible with both types of meaning relations between Sam scratching his cheek and notifying his co-conspirator (25-a).

(25) a. Sam scratching his face means he has a bad hand. *ambiguous, A or B*
b. By scratching his face, Sam means that he had a bad hand. *non-natural (B) only*

However, the syntactic properties of the version in (25-b) render the sentence unambiguous; it can only be felicitously and truthfully used to characterize the case of non-natural meaning in Scenario B. In what follows, I will argue that the (morpho-)syntactic differences between (25-a) and (25-b) – e.g. the difference in the kind of subject (*Sam* vs. *Sam scratching his face*) – track not (non-)natural meaning, but agentivity and eventivity. And, while Grice was right that natural meaning is non-agentive (and stative), that's where the correlation stops. I argue, instead, that Grice's **Passivizability** and **Agentivity** tests are tests for agentivity; and Grice's **Actuality**, **Direct Quotability**, and **Factiness** tests are tests for eventivity.

3.2 Explaining agentivity differences in meaning constructions

We have independent knowledge from the lexical-semantic literature about some semantic properties of verbs in the *cover/amuse* class that I claim *to mean* belongs to (Rappaport Hovav, 2017; Wilson, 2020): in particular, the triadic form asymmetrically entails the unaccusative form.[7]

(26) a. Una covered the screen with the blanket (yesterday). → The blanket covered the screen (yesterday).
b. Karen amused Adam with her antics (yesterday). → Her antics amused Adam (yesterday).
c. By 'Knucklehead,' Ernie meant Peter (yesterday). → 'Knucklehead' meant Peter (yesterday).

For the same reason, the triadic form also entails existential versions of the unaccusative and agentive forms.

(27) a. Una covered the screen with the blanket. → Something/Someone covered the screen (yesterday).
b. Karen amused Adam with her antics. → Something/Someone amused Adam (yesterday).
c. By 'Knucklehead,' Ernie meant Peter. → Something/Someone meant Peter (yesterday).

This is just Grice's **Agentivity** test: the observation that an agentive or triadic form of a meaning construction entails that someone was an agent in that meaning relation.

Crucially, natural meaning cannot be encoded in an agentive or triadic meaning construction, because these constructions mark an agent argument, and natural meaning is agentless. If smoke means fire by virtue of the fact that it is causally related or correlated to smoke in a particular, metaphysical way, that relation is not and cannot be mitigated by any individual: there is no possible agent to specify. Grice's Agentivity test works on meaning constructions that specify an agent, and natural meaning constructions cannot specify an agent, so the Agentivity test is an indirect test for non-natural meaning.

To test this claim, we cannot look for an instance of natural meaning with an agent, but we can find an instance of meaning that is intuitively non-natural (Schiffer, 1972), and test it in an agentless construction (28):[8]

(28) The smoke alarm means fire. ↛ Somebody or other meant fire by the smoke alarm.

The construction in (28) demonstrates that Grice's Agentivity test is actually a test of whether or not a construction specifies an agent, not whether it encodes natural or non-natural meaning. Crucially, natural meaning cannot have an agent, so an agent cannot be specified.

We have observed that agents can be optionally specified for non-natural meaning relations; there is an open question—brought up by Grice himself, in his musing about the traffic light corporation, quoted on the previous page—of whether non-natural meaning requires an agent (a question that is independent of whether or not specification of that agent is required). I think that the answer to this question depends heavily on metaphysical theories about what qualifies as an agent (i.e. a corporation or inventor of a traffic light) that I am not willing to commit to here. But, following the suggestion of two reviewers, I find it plausible to think that stative or habitual

[7] In calculating these entailment relations, we need to hold fixed how we are interpreting them, i.e. we need to control for the eventive or habitual interpretation. To do this, it's best to insert a temporal modifier like *yesterday* into the sentence (to bring out the eventive interpretation).

[8] A reviewer worries that the relationship between a smoke alarm ringing and fire is potentially natural, because the alarm is triggered by 'a certain electrophysical reaction' caused by fire. But this only happens when things go well; as we have all experienced, smoke alarms have a lot of false positives, set off by smoke in the absence of fire, or even by faulty batteries. Natural meaning has no false positives; the signal follows naturally and necessarily from the signaler.

or middle uses (e.g. *Politicians bribe easily*) of all verbs erase a sense of a particular agent without obviously erasing the abstract sense of an agent, and that same tendency is likely at play here, for non-natural meaning relations.

Grice's **Passivizability** test, too, is a test for agentivity, not (non-)naturalness. To make this argument, I'll need to focus for a minute on the difference between the unaccusative form of a meaning construction, as in (29-a), and a passivized version of a triadic form (29-b).

(29) a. 'Knucklehead' meant Peter. *unaccusative form*
b. Peter was meant by 'Knucklehead'. *passivized triadic form*

As we saw in §2, the unaccusative form is one possible syntactic configuration of a meaning construction; in particular, one with a signal subject, and the signal's meaning as a thematic object (of the form 'σ means μ'). We get this form for free from the argument structure of the verb *to mean*. The passive form in (29-b), in contrast, is a passivized form of the triadic configuration *By 'Knucklehead,' someone meant Peter*, i.e. 'μ was meant by σ'.

While they both have two arguments – and, in particular, while they both have the same two arguments – unaccusative and passive forms differ in that the passive form, at least of this sentence, has an implicit agent.[9] This is a familiar claim from the transitive sentence in (30-a) and its passivized version in (30-b).

(30) a. The president made mistakes.
b. Mistakes were made.

The source of the controversy surrounding Nixon's utterance of (30-b) is that it entails someone made the mistakes, by virtue of its relation to the transitive sentence in (30-a) (Cinque, 1988; Collins, 2024). But this property of passive constructions is not a property of unaccusative constructions like (29-a)

To summarize: triadic forms (really, any form that specifies an agent) can be passivized, in which case the agent argument disappears. When this happens, the passive morphology indicates the (syntactic) presence of an implicit agent (which is not a property of the unaccusative form, despite it also lacking a specified agent). Since natural meaning relations can't have agents, they cannot be passivized; in contrast, non-natural meaning relations can have agents, so they can be passivized.

As before, the correlation between agentivity and (non-)naturalness isn't perfect, so we can push on the one place in which they fall apart. (31) is an example of a natural meaning construction failing the passivizability test (as Grice predicts), but (32) is an example of a non-natural meaning construction also failing the passivizability test, in contrast to Grice's predictions.

(31) a. Smoke means fire.
b. *Fire is meant by smoke.

(32) a. The smoke alarm means fire.
b. *Fire is meant by the smoke alarm.

The unaccusative form in (32-a) cannot be passivized because it, too, lacks an agent.[10]

[9]There are constructions whose passive form deletes arguments other than agents; in e.g. *The door was opened* (*by the key*), the targeted argument is an instrument. But in the case of the triadic form, it's the agent that passivization eliminates. Thanks to a reviewer for helpful clarification on this point.

[10]This argument, unfortunately, comes with a big caveat (thanks to a reviewer for extensive and patient discussion here). It is possible that (29-b) is a passive of the unaccusative form, rather than the triadic form. If it is indeed the passive of the triadic form, as I claim above, we would expect the passivized argument to be able to be optionally added via a *by*-phrase (as it is in *Mistakes were made by the president*). But it doesn't appear that this is the case.

(i) a. ??Peter was meant by 'Knucklehead' by Ernie.
b. ?By 'Knucklehead,' Peter was meant by Ernie.

While the position of the *by*-phrase is inexplicably important, as it is even in (29-b), the unacceptability of the sentences in (i) are unexplained in an account that assumes (29-b) is a passive of the triadic form.

The alternative is that (29-b) is in fact the passive of the unaccusative form. After all, other verbs in the *cover* class of verbs can form passives (e.g. *The screen was covered by the blanket*, and, perhaps less acceptable, *?By Una, the screen was covered by the blanket*). If this

3.3 Explaining eventivity differences in meaning constructions

The three remaining linguistic diagnostics from Grice have in common that they test for another property that loosely tracks the (non-)natural meaning distinction: eventivity. The correlation between non-natural meaning and eventivity, however, is slightly more complicated than the one between non-natural meaning and agentivity.

In this subsection, I'll argue that Grice's **Actuality**, **Direct Quotativity**, and **Factiness** diagnostics are actually tests for whether or not a signal argument σ is eventive. This is slightly different than a claim about the eventivity of the meaning relation, but it's correlated. In particular: for natural meaning, the meaning relation is always stative, but the signal argument of this habitual relation can be eventive. For non-natural meaning, the meaning relation can be stative or eventive, and when it is eventive, its signal must be eventive. I'll explain in detail in what follows.

Recall from §2.2.2 that the signal argument σ of meaning relations can denote an eventuality instead of an individual. This is illustrated in (33), repeated from (25).

(33) a. Sam scratching his face means he has a bad hand. *ambiguous, A or B*
b. By scratching his face, Sam means that he had a bad hand. *non-natural (B) only*

In (33-a), the signal subject *Sam scratching his face*, is event-denoting; in (33-b), a triadic form, the signal is encoded in a *by*-phrase, and the subject is now the agent of the meaning relation.

3.3.1 Direct Quotativity is a test for event signals

Grimshaw (2015) observes that *to say* constructions with non-agentive subjects "cannot have quotes in complement position" (p88, exemplified in (34)). It is one of a number of properties she correlates with the difference between agent and non-agent subjects in verbs of saying constructions.

(34) #The sign said, 'The park is closed.'

And while I don't particularly mind (34) – these judgments about orthography are not particularly strong – there is a precedence in the literature for relating agentive subjects with the licensing of direct quotes.

Insofar as we have judgments about the acceptability of quotation marks, I argue here that Grice's **Direct Quotability** test is a diagnostic for whether or not a signal argument σ in a meaning relation is event-denoting (as opposed to individual- or state-denoting). (This is slightly different from Grimshaw's claim about *to say*, as the non-agentivity of a subject is only indirectly correlated with its eventivity). In particular, we can report a meaning μ in quotation marks iff its signal is an event. Intuitively, we can only quote an instance of meaning, not a meaning itself.

Two of Grice's original examples of the Direct Quotability test are in (35) (from (7)).

(35) a. *The recent budget means 'We will have a hard year.' *natural*
b. Those three rings on the bell mean 'The bus is full.' *non-natural*

But these examples do not properly control for the eventive status of the signal argument σ. When we control for eventivity, we see that the Direct Quotability test tracks whether the signal denotes an event, rather than whether the meaning relation is (non-)natural.

The constructions in (36) are instances of natural meaning, and (contrary to Grice's generalization) are perfectly acceptable with their theme or meaning μ in direct quotes. In contrast, the constructions in (37) are instances of non-natural meaning, and (contrary to Grice's generalization) are unacceptable with quotation μ arguments.

---

is right, then we cannot appeal to agentivity to explain the contrast in (31) and (32), because unaccusative forms do not mark agents.

I still believe that Grice's original passivizability test results are best explained by the correlation to agentivity and not (non-)naturalness, and that this is underscored by the data in (32). But I cannot provide clear evidence that (29-b) is a passive of the triadic form. My suspicion is that the ungrammaticality of (i) has to do with other unusual properties of *to mean* (see §4.1) combined with a general oddness of stacking *by*-phrases (cf. *?By Una, the screen was covered by the blanket*), but I cannot advocate for that position more definitively.

(36) **natural meaning** constructions with eventive signals
   a. Smoke billowing means, 'There's a fire!'
   b. Flowers blooming means, 'It's spring!'

(37) **non-natural meaning** constructions with non-eventive signals
   a. *The tricolor flag means, 'Croatia!'
   b. *Flowers in bloom mean, 'It's spring!'[11]

In (36), we know that the natural meaning signals are eventive because the subjects are gerunds. As a result, the μ or theme argument of these natural-meaning relations can appear in quotes. I'll emphasize that the verb *to mean* is nevertheless stative here, denoting a state of signification between an event signal and the meaning theme, and we know this because the verb is in the simple present.

In contrast, in (37), the subjects encoding the signals denote states, not events. As a result, despite the fact that the sentences characterize non-natural meaning relations, the μ or theme arguments cannot appear in quotes.

In sum, Grice characterized the test as a diagnostic of (non-)natural meaning. But it is in fact a test for whether or not the signal σ argument of a meaning relation is event-denoting. Effectively, if the signal is an event, its meaning can appear in quotes. It's possible that Grice missed the mark here because there is a tendency of event-denoting signals to be signals of episodic meaning relations, and episodic meaning relations can only be non-natural meaning relations. But the data in (36) and (37) show that the two properties fall apart, and that Direct Quotativity – insofar as we have intuitions about it – tracks the event status of a signal instead of (non-)naturalness.

3.3.2 Actuality and Factiness are tests for the actuality of eventive signals

If an argument is event-denoting, it may or may not come with an actuality implication, i.e. an implication that the event occurred (in the past, in the actual world).[12] There are at least two ways to condition an actuality implication for an event-denoting argument: 1) include proximal marking in the event-denoting subject; 2) place the event-denoting subject in a construction that is interpreted eventively. I will illustrate each in turn.

The sentences in (38) encode natural meaning relations. In (38-a) the signal σ is an event-denoting DP, the event of lightning striking. Because the verb *to mean* is in the present tense, this meaning relation is most naturally interpreted as stative (see §2.2.1).

(38) a. Lightning striking means a storm is near.
   b. That lightning striking means a storm is near.
   c. Yesterday, lightning striking meant that a storm was near.

The next two sentences represent two ways to lend an actuality implication to the event-denoting subject. In (38-b), this is accomplished using proximal marking (*that* as opposed to a null determiner; see also Acton and Potts 2014). The presence of this proximal marking is enough to make the lightning strike be interpreted as actual: the speaker is committed to it having occurred in the actual world (and, in fact, in the immediate, proximal past of the actual world). Crucially for us, the main verb *to mean* is still interpreted statively: the meaning relation is still interpreted as non-contingent, holding over time. It's just that (38-b) expresses that the very recent and actual lightning strike is an instantiation of that habitual meaning, especially in relation to the meaning. Since the signal σ is actual, then, assuming that the speaker is an authority who is speaking truthfully, we can infer the actuality or truth of the attributed meaning μ, i.e. that a storm is near.

In (38-c), we achieve an actuality implication on the event-denoting argument in a different way: by triggering a lifetime effect. When stative verbs occur with event modifiers, we get an implicature that the relation no longer

[11] Thanks to a reviewer for pointing out the usefulness of this contrast with (36-b).

[12] It is more typical to refer to this as an actuality entailment, not an actuality implication. But because I am relying on the claim that this actuality arises in complicated ways that involve cooperation between aspectual and pragmatic properties, I will relegate it to the less committal property of an implication, to avoid controversy on that particular front.

holds; this is true, too, for (38-c), which has the odd effect of suggesting that this meaning relation between lightning and storms no longer holds. The same goes for the similar sentence in (39).

(39) Smoke yesterday meant that the cardinals elected a new pope.

If we are instead looking at a meaning relation that is eventive, or contingent – and particular, one that happened in the past – then we get an actuality implication for the event-denoting subject. In particular, the claim "A signal σ, in the past, indicated that μ" carries with it the implication that σ actually took place. Note that this claim about actuality is only relevant to event-denoting σ arguments, not individual-denoting ones. And again, if σ actually occurred, assuming that the speaker is authoritative and telling the truth, we can infer the actuality of the meaning μ the speaker attributes to it.

In sum, there is an indirect correlation between whether a signal σ argument is event-denoting, and whether the meaning relation in which it participates is eventive. In other words, if a meaning construction has an eventive signal σ with an actuality implication, it is likely to be the argument of a meaning relation that is being interpreted eventively. If a meaning relation is interpreted eventively, it is likely to be a non-natural meaning relation (i.e. to be contingent).[13]

Two additional notes here: one, this is an asymmetrical entailment, i.e. non-natural meaning relations don't *need* to be eventive, and thereby don't need to have actual eventive arguments. And two, there's a way to achieve an actuality implication that doesn't involve an eventive meaning relation: by adding proximal marking to the event-denoting subject. This will allow us to test the present claim that Grice's diagnostics are only indirect tests of (non-)natural meaning.

Some of Grice's original examples are in (40).

(40) a. Those spots mean measles, #but he hasn't got measles. *natural meaning*
b. That bell ring meant the bus is full, but it isn't in fact full. *non-natural meaning*

Compare (41) to Grice's original test in (40-a): the signal argument in (40-a) contains proximal marking, which prompts an actuality implication for the event associated with the subject, resulting in Grice's original intuition that the sentence entails that the person in question has spots (and therefore, given the stativity of the meaning relation, that he has measles).

(41) Ring-shaped spots mean measles, but no one here has {ring-shaped spots/measles}.

In contrast, the signal argument in (41) lacks proximal marking, and thereby an actuality implication. As a result, there is no evidence of an actual signal, and so the speaker is no longer committed to anyone having measles.

In (42), we can test how event actuality relates to non-natural meaning. Instead of talking about bus bells, which are typically rung, I'll switch to the non-natural meaning scenario involving Sam's poker game, the one in which he scratches his face to signal to a co-conspirator that he has a bad hand. The version in (42-a) is a parallel of the Grice example in (40-b): it's a non-natural meaning construction that passes the non-actuality test, i.e. that lacks an actuality implication.

(42) a. Sam scratching his face means he has a bad hand, but he doesn't have a bad hand.
b. Sam scratching his face meant he had a bad hand, #but he doesn't have a bad hand.

However, one thing contributing to the non-actuality status of (42-a) is the tense on the verb: it is in present tense, which means it's most naturally interpreted statively. But when we change the tense to past tense in (42-b), we switch to an eventive interpretation of the meaning relation. And, in turn, we get an actuality implication for the meaning argument, which means that we have a non-natural meaning construction that fails the non-actuality test, contrary to Grice's generalization.

[13] Anand et al. (2017) has made a similar point for a group of terms they dub 'communicative predicates', i.e. that they behave differently from other sentence-embedding verbs (doxastic, emotive, inferential) in that they allow for certain inanimate objects in the theme or instrument role, and in that they are never factive (on the μ argument).

We can run the same arguments for Grice's **Factiness** diagnostic, as it, too, is a test of event actuality; only event-actual signal arguments can be correctly paraphrased as facts of the matter. This is demonstrated by the sentences in (43), in which both natural and non-natural meaning relations pattern together on the Factiness test. This is because both main verbs are in the past tense, and thereby both constructions carry an actuality implicature for the event-denoting signal argument.

(43) a. The fact that he had those spots meant that he had measles.
b. The fact that Sam scratched his face meant that he had a good hand.

To sum up this section: Grice's diagnostics are in fact tests of the agentivity and eventivity of the meaning relation. Because natural meaning is necessarily agentless, it cannot be encoded in an agentive or triadic form, and therefore doesn't pass the Passivizability or Agentivity tests. Because natural meaning is generally stative, natural meaning constructions are perhaps less likely to have eventive signal arguments, or at least eventive signal arguments with an actuality implication. But they can, and when natural meaning constructions do have eventive signal arguments, they pattern just as non-natural meaning constructions with eventive signal arguments do.

The goal of the first half of this paper has been to argue that Grice's linguistic diagnostics, instead of diagnosing the (non-)natural meaning distinction, test for two loosely correlated properties (agentivity and eventivity). In the rest of this paper, I address the question of polysemy: namely, whether we should take all of this to be evidence of two distinct verbs *to mean*, or as normal variants of a single verb.

4 A question of homophony or polysemy

It is relatively easy to read Grice (1957) and as arguing that the import of his linguistic diagnostics is that the word *to mean* is ambiguous between *meaning*$_N$ and *meaning*$_{NN}$ (his shorthands for natural and non-natural meaning, respectively), but it's possible he intended something slightly different. It certainly seems that Grice grows more circumspect about what these two senses amount to for the semantics of the verb *to mean* in 1989 (p291-2):

> "Is this double use of the word 'mean' just like the double use of the word 'vice' to refer sometimes to something approximating to a sin and sometimes to a certain sort of instrument used by carpenters? [...] On general grounds of economy, I am inclined to think that if one can avoid saying that the word so-and-so has this sense, that sense, and the other sense, or this meaning and another meaning, **if one can allow them to be variants under a single principle, that is the desirable thing to do** [emphasis mine]: don't multiply senses beyond necessity."

In this section, I want to emphasize how standard the meaning alternations between (non-)agentivity and (non-)eventivity are for verbs, as discussed in §2.2.1, and what that means for the meaning of *to mean*. In particular, I argue that there is no reason to postulate two distinct verbs *to mean* in English, assuming it's less theoretically desirable to do so, because we already have mechanisms for accounting for these alternations in our lexical and event semantics. However, on the flip side, there is nothing intrinsic about the verb *to mean* that requires it be agentive, or non-agentive, or underspecified with respect to agentivity (or eventive, or stative, or underspecified with respect to stativity). There are very similar words in English that are not underspecified in these respects, or at least behave differently from *to mean* in these respects (§4.1); and there are synonymous or near-synonymous words in other languages that lexicalize these requirements (§4.2). So, in summary, we can derive these different interpretations of the meaning relation without postulating two different *to mean* verbs, or even that the verb *to mean* is polysemous.

There are a few different ways to conceptualize two string-identical words with distinct meanings. One is as a circumstance of accidental homophony, cases where two unrelated meanings happen to have a string-identical label (as in *bat* in (44-b)). This is surely not what's happening with Grice's *meaning*$_N$ and *meaning*$_{NN}$, whose meanings are clearly related. This leaves the question of whether or not *meaning*$_N$ and *meaning*$_{NN}$ are polysemous, i.e. have distinct meanings that are systematically related to one another in a predictable way.

There are recent compelling arguments that polysemy is not a coherent notion independent from underspecification (Brody and Feiman, 2023). But, traditionally, at least one type of polysemy is diagnosable by

copredication. Specifically, certain polysemous words like *lunch* (polysemous between a meal's contents and its runtime) to be simultaneously able to occur with a verbal predicate that selects for food (e.g. *be delicious*) and one that selects for events (e.g. *last a long time*), as in (44-a). (This is in contrast with a homophonous DP, like *bat*, as shown in (44-b).)

(44) a. Lunch was delicious and lasted a long time.
b. #A bat is a mammal and a piece of sporting equipment.

If we wanted to pursue a traditional polysemy diagnosis for *to mean*, we might flip copredication on its head: perhaps we can diagnose polysemy in a verb if it can simultaneously predicate two conjoined arguments that work with distinct senses of the verb. This is tested in for the two different possible subject arguments of the verb *to cover* (45-a), and the two different possible subject arguments of the verb *to mean* (the signal σ and the agent x, in (45-b)).

(45) a. #The blanket and Una covered the screen.
b. #'Knucklehead' and Ernie meant Peter.

I don't believe (45-a) can receive the intended interpretation, namely that Una covered the screen with a blanket. And I don't believe (45-b) can receive the interpretation that Ernie used 'Knucklehead' to mean Peter. This means that the verb *to mean* isn't polysemous in at least one sense of the term.

This suggests something that won't be surprising to lexical semanticists: we don't want to think of verbs that exhibit distinct argument realization options as instances of distinct verbs (either homophonously or polysemously distinct): we should instead think of them as a single verb, with any perceived differences in meaning conditioned all and only by the argument alternations that distinguish them. In the next two subsections, I provide two other arguments in favor of this conclusion, namely that the verb *to mean* encodes a single meaning relation, but it is idiosyncratic (relative to other verbs in English, or to meaning verbs in other languages) in what sorts of arguments it can relate, and how.

4.1 Agentivity and eventivity in other English verbs

There are several English verbs that pattern like *to mean* in exhibiting the argument alternation pattern in (14) that is unique to the *cover/amuse* class of verbs (Rappaport Hovav, 2024). They are, as far as I can tell: *to denote; to designate; to indicate; to mean*; and *to refer to*.[14]

(46) a. 'Knucklehead' {denotes/designates/indicates/means/refers to} Peter.
b. Ernie {denotes/designates/indicates/means/refers to} Peter.
c. By 'Knucklehead,' Ernie {denotes/designates/indicates/means/refers to} Peter.

There are some plausibly semantically similar verbs that don't exhibit this argument alternation, which is what we would expect, as we have independent reason to believe that argument structure is not completely dependent on meaning. There are some synonyms of *to mean* that take on a different meaning in the absence of the signal argument, as demonstrated in (47).

(47) #Ernie {represents/says/signals/stands for} Peter.

And other plausibly semantically similar verbs are ungrammatical in the absence of the signal argument.

(48) *Ernie {conveys/expresses/reads as} Peter.

That said, because these verbs all have optional agent arguments (they all allow for sentences of the form *'Knucklehead' [verb] Peter*), we expect at least Grice's agentivity tests to work on them. And they do:

(49) **Agentivity**

[14] A reviewer reports being unable to use the verbs *to indicate* and *to refer* in simple present (in contrast to the author's judgments reported here), requiring instead the present progressive in their versions of the sentences in (46). For this reviewer, this is a strong indication that these verbs are agentive in these forms.

a. Smoke indicates fire. ↛ Somebody used smoke to indicate that there is fire. *natural*
b. 'Knucklehead' indicates Peter.
↛ Somebody used 'Knucklehead' to indicate Peter. *non-agentive non-natural*
c. By 'Knucklehead,' Ernie indicated Peter.
→ Somebody used 'Knucklehead' to indicate Peter. *agentive non-natural*

(50) **Passivizability**

a. ??What was indicated by smoke was fire. *natural*
b. What was indicated by 'Knucklehead' was Peter. *non-natural*

There is another dimension along which these verbs vary: while *to mean* can be stative or eventive, depending on the context and other factors, some of these verbs are lexically more stative, and some lexically more eventive. The verb *to denote*, on the one hand, is quite stative: it can't describe a particular event, and as a result can't e.g. head an imperative (cf. **Denote Peter* (*when you use the word 'Knucklehead'*)*!*). On the other hand, the verb *to refer to* is quite eventive: it can only describe an eventive relation, and can, in contrast, head an imperative (e.g *Refer to Peter* (*when you use the word 'Knucklehead'*)*!*). It seems closer to *to say* than to *to mean* in this respect (see §2.2.1).

Given this variation in lexical aspect, we would expect Grice's tests for eventivity to generalize less well to the full class of signifying verbs, albeit in predictable ways. For instance, we'd expect stative verbs like *to denote* to have a hard time with eventive interpretations across the board, regardless of the (non-)naturalness of the meaning relation. And this is what we see:

(51) a. *The bell rings denote, 'The bus is full.' (*non-natural meaning fails Direct Quotability*)
b. The spots denote measles, but no one here has spots. (*natural meaning fails Actuality*)

And we would expect eventive verbs, like *to refer to*, to have a hard time across the board with stative interpretations, and therefore any sort of natural meaning relation. And this is what we see:

(52) a. *Smoke refers to fire.
b. *Those spots refer to measles.

In sum, there's evidence that *to mean* patterns like other signifying or communicative verbs in English in several respects: its syntactic alternations (in particular, its ability to occur with an agentive and/or signal argument); and its aspectual polysemy (in particular, its ability to be interpreted statively or eventively, depending on the nature of its arguments, and the morphosyntactic properties of the construction).[15]

What does this mean for Grice's original claims about meaning? It seems likely that there are no direct, perfect linguistic reflexes of the (non-)natural meaning distinction. But the good news is that we can instead reconceive of the contribution of Grice (1957) as providing interesting linguistic evidence for the correlation between non-natural meaning and (syntactic) agentivity: namely that natural meaning relations cannot specify an agent argument, while non-natural meaning relations can. Because this claim is clearly closely related to Grice's claim that non-natural meaning is distinguished by the presence of a psychological agent – the instigator of a psychologically mediated communicative act – it is still quite in line with the Gricean program. Natural meaning is meaning that always lacks an agent, and this sort of meaning is generally expressed in stative terms.

In the final section of this paper, I argue that the linguistic distinctions argued for above are not just the right ones, but are also not specific to English. In other words, that when a natural language employs more than one word for *to mean*, the lexical distinction tracks agentivity and eventivity, not (non-)natural meaning.

[15] A reviewer points out that, morphosyntactically, *to mean* behaves quite in line with a semantically related group of verbs: *to explain*, *to justify*, and *to prove*. All of these verbs seem to have baseline stative interpretations, but can also be used agentively.

(i) a. His background explains/justifies his behavior.
b. Daniel explained/justified his behavior (by referring to his background).

4.2 Agentivity and eventivity in other languages

There are, presumably, other languages that can use the same verb *to mean* for both natural and non-natural meaning relations, and agentive/non-agentive meaning relations. But there are lots of languages that use distinct words or phrases for meaning relations, and in all of these languages that I've seen, the lexical difference tracks the (non-)agentive distinction, not the (non-)natural distinction.

The goal of this section is to argue that, for these languages, lexical distinctions in signifying verbs actually track agentivity (and, to a lesser extent, eventivity), not (non-)naturalness. In eliciting data for this section, I consulted with linguist and philosopher consultants on their native language. Because it's quite common for languages to employ multiple synonyms or near-synonyms for *to mean*, just as English does (cf. §4.1). For this reason, I focused on languages whose consultants used two different signifying verbs in initial elicitations of two canonical natural and non-natural meaning examples, e.g. 'Smoke means fire' and '[word] means fire.'

(53) is an example of (non-)natural meaning both encoded in the same morpheme in Korean; the difference is that *ttus* is a nominal, instead of a verb.

(53) a. σ-uy **ttus**-un μ-ta.
σ-GEN meaning-TOP μ-be
'The meaning of σ is μ.'
b. x-hante σ-**ttus**-un μ-ta.
x-DAT σ-GEN meaning-TOP μ-be
'σ means μ to x.' *Korean*

Greek has (at least) two different root words corresponding to *to mean*:

(54) a. O kapnos **simeni** fotja.
the smoke mean.PRS.IMPF.2SG fire
'Smoke means fire.'
b. O Nikos **enoi** fotja me tin leksi 'fuego'
the N mean-PRS.IMPF.3SG fire with the word f
'(By 'fuego',) Nikos means fire.' *Greek*

The verb in (54-a) is what was elicited for natural meaning, and the verb in (54-b) the one for non-natural meaning. But, crucially, the distinction between the two is that the former doesn't have an agent argument, while the latter does. This is evident in the unaccusative but non-natural meaning relation in (55).

(55) I leksi fotiá **simeni** fotja
the word 'fotiá' mean-PRS.IMPF.3SG fire
'The word 'fotiá' means fire.' *Greek*

This construction uses the unaccusative, non-agentive verb *simeni*, rather than the agentive verb *enoi* previously used for the non-natural meaning relation in (54-b).[16]

Spanish is an exemplar of Romance languages. It has a two-place verb *significar*, which does not allow for an agent, i.e. is only unaccusative.[17]

(56) a. Fuego significa corre.
fire mean-3.SG run
'Fire means run!'
b. (?A Laura), coffee significa café.
to L coffee mean-3SG coffee
"Coffee' means coffee.' *Spanish*

[16]Speakers of several other languages – Estonian, Icelandic, and Romanian – have reported to me the same configuration.

[17]It is possible to associate the verb *significar* with an experiencer, as in *Para mi* ('For me')…. But this is different than an agent argument, syntactically and in other ways. I'll return to discuss experiencer arguments in §5.

It also has a second phrase, 'wants to say,' that requires a syntactic agent. (57) exemplifies the Spanish version of the English triadic construction.

(57) Con 'coffee', Isa quiere decir café.
with coffee Isa want-3SG to-say coffee
'By 'coffee,' Isa means coffee.' *Spanish*

It is *possible* for a canonically natural meaning relation to be encoded in this 'wants to say' construction, but forcing a non-canonical agent into the agent position in this construction results in a marked, non-natural interpretation.

(58) Fuego quiere decir corre.
fire want-3SG to-say run
'Fire means run!' (lit. 'Fire wants to say run.') *Spanish*

The use of the volitional *querer decir* in (57) is only appropriate in a subset of conditions that (56-a) can be used: it only makes sense in a context in which there's a sign with a fire symbol on it, i.e. something that's personifiable. This is much like the effect we get in a forced triadic version of a natural meaning relation in English, e.g. *By burning, the fire meant smoke*, which invokes a sentient fire (à la *Howl's Moving Castle*).

Interestingly, there are lots of other, unrelated languages that employ some complex 'wants to say' constructions for the agentive version of *to mean*. Avatime, a Niger-Congo language spoken in Ghana, can use the copula to express habitual meaning relations, but requires a verb of saying to express eventive meaning relations (and this distinction cross-cuts the (non-)natural meaning distinction).[18]

(59) a. ki-mimi-e ki-nu rice (ki me)
CL-rice-DEF CL-be rice to 1SG
"Kimmie' means rice (to me).'
b. ki-mimi-e me-do si ki-nu rice
CL-rice-DEF 1SG-tell say CL-be rice
"Kimmie' I say to mean rice." *Avatime*

Both versions of 'to mean' in Uyghur (a Turkic language spoken in China) involve the verb 'to say,' but the language distinguishes between stative and eventive meaning relations via other means: an eventive (non-natural) meaning relation in (60-a) involves a 'meaning' nominal, amounting to something like 'The word's meaning when one says 'bodek' is fat'.[19]

(60) a. bodek dé-gen söz-ning meni-si sémiz
fat say-PTCP.PST word-GEN meaning-3POSS fat
'The word 'bodek' means fat." (lit. 'The word's meaning when one says 'bodek' is fat.')
b. is de-gen-lik ot de-gen-lik
smoke say-PTCP.PST-C fire say-PTCP.PST-C
'Smoke is fire.' (lit. 'So-called smoke is equal to so-called fire.') *Uyghur*

While stative meaning relations (e.g. (60-b)) involve a subjectless use of 'say' as an impersonal marker that is restricted to stative or generic contexts (see Major (2024) for more detail on this construction).

In sum, lots of languages employ a variety of different words or strategies for specifying meaning relations. But the preliminary evidence presented here suggests that, when a language has more than one word for 'to mean' that appear to alternate with the (non-)natural meaning distinction, the distinction is really tracking something like agentivity (in the case of Greek and Spanish) or eventivity (in the case of Uyghur). This provides additional evidence that the linguistic reflexes we see of the (non-)natural meaning distinction are really just indirect symptoms of other semantic distinctions that we have independent evidence of: in particular, the presence or absence of an agent argument; and whether or not the relation is thought to hold eventively or habitually.

[18]Thanks to Divine Agyepong and Vincent Azafokpe via Travis Major for the Avatime data.
[19]Thanks to Gulnar Eziz and Akbar Amat via Travis Major for the Uyghur data.

5 Conclusions and extensions

Recall that Grice (1989) admitted that "if one can avoid saying that the word so-and-so has this sense, that sense, and the other sense, or this meaning and another meaning, if one can allow them to be variants under a single principle, that is the desirable thing to do" (p291-2). The goal of this paper is to do this for the verb *to mean*, by bringing more linguistic rigor to Grice's original five tests separating out natural and non-natural meaning relations. The claim here is that Grice's tests targeted agentivity and eventivity, which i) are independently known to be marked morphosyntactically; and ii) correlate to varying degrees with (non-)natural meaning. So we can maintain a single sense of the verb *to mean*, and attribute all apparent differences in linguistic behavior to the presence or absence of an agent, and the eventivity of the verb in a given construction, just as we do for other verbs.

If this is right, then the contribution of the present paper is to clean up Grice's linguistic claims and to precisify his conclusions. In particular, there are no observable linguistic reflexes of the (non-)natural meaning distinction, but there are linguistic reflexes of semantic correlates of the (non-)natural meaning distinction: namely the presence or absence of an agentive argument (what I've called 'agentivity') and whether a meaning relation can hold eventively, or contingently (what I've called 'eventivity'). If we take the reasonable position that natural meaning, just intrinsically, cannot be specified with an agent argument and cannot be interpreted eventively, I argue that Grice's linguistic observations are accounted for.

But while Grice was wrong about the linguistic manifestation of (non-)natural meaning, he was still right about a lot of things. The Gricean project, as I interpret it, was to urge people interested in natural language meaning to consider the signaler, in addition to the signal and the meaning signaled. In Grice (1975) and other work, this has quite famously and, in my mind, successfully resulted in a project in which philosophers of language and semanticists consider the agent of a communicative act, and their goals. From that perspective, the main contribution of Grice (1957) was to argue that the principal difference between natural meaning and non-natural meaning is that the latter, but not the former, involves a psychological agent. Now that we know what to look for, it is this key factor that we observe in the linguistic tests Grice put forward.

This obviates the need for any linguistic distinction between natural and non-natural meaning. Which is to say, we don't need to think of the English verb *to mean* as homophonous or polysemous or having more than one sense, at least in contrast to any other verb of communication or signification, assuming this is theoretically less desirable. Verbs like the English *to mean* can be interpreted eventively or not; and they can specify an agent or not. And when we find languages that lexically differentiate between senses of meaning, we find that these lexical distinctions track agentivity and eventivity, not (non-)naturalness of meaning.

I take it then, to mean that the Gricean program can remain intact, or at least that the details of Grice (1957) do not impact the overall goal of associating conventional meaning with a psychological agent. That said, it's important not to treat the verb *to mean* as anything special in any other sense except for maybe a narrow syntactic sense, given its quirky argument alternation patterns. And it's always worthwhile examining other languages when drawing conclusions from linguistic behavior to the state of the world.

I'll end by speculating about one other way in which this paper might be useful: I have specified three semantic arguments of *to mean*: the signaler x, the signal σ, and the meaning μ. But meaning relations can also specify an experiencer, introduced by the preposition *to* (61-c).[20]

(61) a. 'Knucklehead' means Peter.
b. By 'Knucklehead,' Ernie means Peter.
c. 'Knucklehead' means Peter to Ernie.

This means that there are a total of four syntactic arguments that can associate with the meaning relation: i) a

[20]Crucially, the experiencer can but need not be the agent: while the sentence *By 'Knucklehead,' Ernie meant Peter* entails that Ernie used the term 'Knucklehead' in a speech act, the sentence *'Knucklehead' meant Peter to Ernie* is compatible with Ernie being a hearer, or in the audience, of a speech act.

meaning μ; ii) an agent x; iii) a signal σ; and iv) an experiencer y. This is schematized in (62).[21]

(62) a. σ meant μ. *unaccusative*
b. x meant μ. *agentive*
c. By σ, x meant μ. *triadic*
d. σ meant μ to y. *experiencer*

Of these four arguments, fully three can be left implicit: the only argument all of these forms have in common is the meaning, μ argument. As a result, the verb *to mean* is a new playground for those of us interested in the semantics of implicit arguments. And, further, we might be able to learn a bit more about the meaning relation by playing around with some of these implicit arguments.

Just like with other predicates with implicit arguments, the relation between a signal and its signified can vary. And it turns out we can use this variation to diagnose subtle differences between types of meaning. I consider three types of terms: common terms, whose meaning is understood by the general population; specialist terms, whose meaning is understood by a strict subset of the general population; and what I dub 'hijacked terms,' which have one meaning among the general population, and a distinct meaning within a specialist community. I take David Chalmer's *zombie* to be an example of the latter, but I don't mean to suggest that a hijacked term can only come from the general population to the specialist community; the general population can hijack specialist terms, too, as exemplified by the common (mis-)understanding of the word *OCD*.

There seem to be real linguistic reflexes of this three-way contrast in types of meaning. We can test the difference on implicit agent arguments as well as implicit experiencer arguments. In (63)-(65), I do this for implicit agent arguments.

(63) **common terms**
A: *blah blah blah* car.
B: #What does 'car' mean?
B': #What do you mean by 'car'?

(64) **specialist terms** (where B is a non-specialist)
A: *blah blah blah* morpheme.
B: What does 'morpheme' mean?
B': What do you mean by 'morpheme'?

(65) **hijacked terms** (where B is a non-specialist)
A: *blah blah blah* zombie.
B: #What does 'zombie' mean?
B': What do you mean by 'zombie'?

In the case of a common word, like *car*, there is no variation in understanding from the general population to any sort of specialist population. It is therefore weird to ask what 'car' means (cf. (63)B), but it is also weird to ask what someone means by 'car' (as in (63)B') as the question presupposes variation. In the case of specialist terms like 'morpheme,' both questions are acceptable, because the non-specialist neither knows what the term means, nor knows how a specialist intends the term. But hijacked terms exemplify a third category, one in which a non-specialist can ask about the meaning of the term relative to a specified agent (a specialist), but cannot felicitiously ask about the meaning of the term writ large.

In (66)-(68), I do this for implicit experiencer arguments, based on the empirical patterns from the predicates of personal taste literature (e.g. Stephenson, 2007), with the same result.

(66) **common terms**
A: 'Car' means X.
B: No, it means Y!...

[21]The signal argument σ is considered part of the main argument structure because it can be introduced without a preposition in the unaccusative form, e.g. in (61-a). This is not true of the experiencer, so we consider that a non-core argument, or an adjunct.

B: ...You're wrong.

(67) **specialist terms** (where *B* is a non-specialist)
A: 'Morpheme' means X.
B: #No, it means Y!...
B: #...You're wrong.

(68) **hijacked terms** (where *B* is a non-specialist)
A: 'Zombie' means X.
B: No, it means Y!...
B: #...You're wrong.

We can replicate something like faultless disagreement for meaning relations, but only for hijacked terms. This suggests that, when the meaning of words varies from population to (sub-)population, speakers are aware of this variation, and employ it in their semantic knowledge. If this is right, we should think about modeling this information, and we should consider what role it plays in cases of misunderstandings and semantic disagreements.

In sum, I believe that Grice's notion(s) of meaning remain intact, but that there is a bit more work to be done in critically examining the (semantic) argument structure of the verb *to mean* and its synonyms in English and other languages. Doing so might result in a better understanding of the meaning relation; of implicit arguments; and of metasemantic knowledge.

References

Acton, E. and Potts, C. (2014). That straight talk: Sarah Palin and the sociolinguistics of demonstratives. *Journal of Sociolinguistics*, 18:3–31.

Anand, P., Grimshaw, J., and Hacquard, V. (2017). Sentence embedding predicates, factivity, and subjects. In Condoravdi, C., editor, *Lauri Karttunen FestSchrift*. CSLI Publications.

Baker, M. and Vinokurova, N. (2009). On agent nominalizations and why they are not like event nominalizations. *Language*, 85:517–556.

Brody, G. and Feiman, R. (2023). Polysemy does not exist, at least not in the relevant sense. *Mind & Language*, 39:179–200.

Cinque, G. (1988). On *si* constructions and the theory of ARB. *Linguistic Inquiry*, 19:521–581.

Collins, C. (2024). *Principles of argument structure: a Merge-based approach*. MIT Press.

Comrie, B. (1976). The syntax of action nominals. *Lingua*, 40:177–201.

Dowty, D. (1979). *Word Meaning and Montague Grammar*. Reidel.

Ginzburg, J. and Sag, I. (2001). *Interrogative investigations: the form, meaning, and use of English interrogatives*. CSLI Stanford.

Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66:377 – 388.

Grice, H. P. (1975). Logic and conversation. In Cole, P. and Morgan, J., editors, *Syntax and Semantics, 3: Speech Acts*, pages 41–58. Academic Press.

Grice, H. P. (1989). Meaning revisited. In Grice, H., editor, *Studies in the Ways of Words*, pages 283–303. Harvard.

Grimshaw, J. (2015). The light verbs 'say' and 'say'. In *Structures in the mind: essays on language, music, and cognition in honor of Ray Jackendoff*, pages 79–99. MIT Press.

Levin, B. (1993). *English verb classes and alternations: a preliminary investigation*. University of Chicago.

Levin, B. and Rappaport-Hovav, M. (1995). *Unaccusativity: at the syntax-lexical semantics interface*. MIT Press.

Major, T. (2024). Re-analyzing 'say'-complementation: implications for case theory and beyond. *Natural Language and Linguistic Theory*, 42:1125–1190.

Moens, M. and Steedman, M. (1988). Temporal ontology and temporal reference. *Computational Linguistics*, 14:15—28.

Neale, S. (2022). Determinations of meaning. In Lepore, E. and Sosa, D., editors, *Oxford Studies in Philosophy of Language vol. 2*. Oxford University Press.

Rappaport Hovav, M. (2017). Grammatically relevant ontological categories underlie manner/result complementarity. In *Proceedings of IATL 2016*, volume 86, pages 77–98. MIT Working Papers in Linguistics.

Rappaport Hovav, M. (2024). Uncovering the scale: On the interaction between the semantics of roots and functional structure. In *Proceedings of the West Coast Conference on Formal Linguistics*, volume 39, pages 371–381. Cascadilla Press.

Schiffer, S. (1972). *Meaning*. Oxford.

Stephenson, T. (2007). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30:487–525.

Sturman, B. (2021). *Subclausal quotation intonation*. PhD Thesis, University of California, Los Angeles.

Vendler, Z. (1957). Verbs and times. *Philosophical review*, 56:143–160.

Wilson, M. (2020). The reversible core of object experiencer, location, and *govern*-type verbs. In *Proceedings of NELS*, volume 49, pages 285–294.