

Agentive meaning

In his foundational 1957 paper, Grice introduced a series of linguistic diagnostics for a familiar or at least intuitive distinction between natural and non-natural meaning. He argued that the verb *to mean* is ambiguous between uses (like *Smoke means fire*) for which the relation specified by the verb is not in the province of semantics (natural meaning); and uses (like *'Fuego' means 'fire'*), in which the verb relation is (non-natural meaning). The essential argument in Grice (1957), sourced quite heavily from his linguistic claims, is that non-natural meaning amounts to psychologically-mediated actions by individuals; i.e. by psychological agents. This claim was the catalyst for the entire Gricean project, which centers speakers and signallers, rather than language, in the communicative act.

The paper has also, however, spurred a great deal of debate about the role of psychological agents in non-natural meaning; the nature of the meaning relation; and the relationship between the (non-)natural meaning distinction and other linguistic phenomena present in meaning constructions.

The goal of the present paper is to re-examine Grice's linguistic argumentation in favor of the (non-)natural meaning distinction with the lens of contemporary linguistic theory. In doing so, I will argue that Grice's linguistic diagnostics are not in fact tests for the (non-)natural meaning distinction, but that they are instead tests for two closely related, independently-recognized linguistic properties: agentivity (whether or not an eventuality has an agent, in the syntactic sense) and episodicity (whether an eventuality argument is a single event, or a habitual state or series of events). The two sets of distinctions fall apart in critical cases for which Grice's interpretation of the theory makes the wrong predictions. The resulting product is a new, more controlled perspective on the classic data, resulting in a more nuanced and, I argue tentatively, more generalizable claim.

I end by demonstrating the pervasiveness of agentivity and episodicity distinctions – but not the (non-)natural meaning distinction – in other verbs in English, and in other languages.

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 **signaller** (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

In §2 I will assign each of these a theta-role and argument status in the linguistic meaning relation, but for now it's sufficient to understand their (potential) role in a meaning report.

The discussion in Grice (1957) takes for granted that we have intuitions about which meaning relations are (non-)natural; the goal of that paper is to provide linguistic diagnostics to underpin the distinction, and to champion Grice's protracted argument that a principle difference between the two is the involvement of an agent, in the psychological or philosophical sense.

The sentences in (2) and (3) are exemplifications of the intuitive (non-)natural meaning distinction, largely lifted 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.

Effectively, 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 you are coming to the table with a distinction between (2) and (3), and he goes on to highlight what he takes to be linguistic reflexes of this distinction.

In his defense of the (non-)natural meaning distinction, Grice presents five linguistic diagnostics.

- a. **Passivizability:** A sentence of the form ‘ σ means μ ’ which conveys a natural meaning relation, like those in (2), does not entail a sentence of the form ‘What was meant by s is σ ’. (This is not true of the sentences that convey non-natural meaning relations, like those in (3).)
 - (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 signallers 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).

- b. **Agentivity:** A sentence of the form ‘ σ means μ ’ which conveys a natural meaning relation, like those in (2), does not entail a sentence of the form ‘Somebody or other meant μ by σ ’. (This is not true of sentences that convey non-natural meaning relations, like those in (3).)
 - (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.

- c. **Actuality:** In cases like (2), we can infer that the speaker is committed to μ being actual; in cases like (3), we cannot.
 - (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 term ‘factivity’ here (cf. Neale, 2022). First, the term 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 phenomenon of factivity – as entailment to a proposition – 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)], ‘[s] meant that p ’ and ‘[s] 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 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

¹I will return to discuss this potential kinship in §5.

the syntactically nominal signifying argument *measles* into a proposition. I do not have such a theory, and I will argue in §4 that these arguments are best analyzed as denoting eventualities, and so am opting for a more non-committal term here.

d. **Direct Quotability:** Sentences like (2) that encode natural meaning relations cannot be paraphrased using quotes for the μ argument; sentences like (3) that encode non-natural meaning relations can.

- (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.

e. **Factiness:** In sentences like (2) encoding natural meaning relations, the signal σ can be truth-preservingly paraphrased as a fact; in sentences like (3), which encode non-natural meaning relations, it cannot.

- (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.

It’s worthwhile mentioning the important if short-lived cottage industry on the semantics of factiness, or the distribution of *the fact that* clausal marking, has had since Grice (1957): in particular, Ginzburg and Sag (2001). These judgments are similarly subtle, 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 Quotability 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?	X	✓
Agentivity (b)	(5)	Does ‘ σ means μ ’ entail ‘Somebody means σ ’?	X	✓
Actuality (c)	(6)	Is the speaker committed to μ being actual?	✓	X
Direct Quotability (d)	(7)	Can the signified σ argument occur in quotation marks?	X	✓
Factiness (e)	(8)	Can the signifier s be paraphrased with ‘the fact that’?	✓	X

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 episodicity (whether the eventuality argument associated with the meaning relation is interpreted as episodic or habitual). The latter will also get tied up a bit in tests for whether a signal σ denotes an actual event.

I will argue that, while natural meaning is necessarily non-agentive and habitual, non-natural meaning can come in any variety, so there are meaning constructions for which Grice’s distinction and mine fall apart. I will also argue that there are tight (morpho-)syntactic reflexes of agentivity and episodicity, so we can witness the syntactic properties of a meaning construction conditioning whether it is interpreted as natural or non-natural meaning. I will illustrate this latter point as a preview.

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 Sam notifying his co-conspirator, e.g. (9-a).

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

However, the syntactic properties of the version in (9-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 (9-a) and (9-b) – e.g. the difference in the kind of subject (*Sam* vs. *Sam scratching his face*) and the tense/aspect on the verb – track not (non-)natural meaning, but agentivity and episodicity. And that these (morpho-)syntactic properties are independently associated with agentivity and episodicity across English, and cross-linguistically.

It is easy for a linguist like me to read Grice (1957) and interpret him as arguing that the import of his linguistic diagnostics is that the word *to mean* is ambiguous or homophonous between *meaning_N* (his shorthand for natural meaning) and *meaning_{NN}* (his shorthand for non-natural meaning), 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.”

My goal here is not to argue that the verb *to mean* doesn't have two senses; it clearly seems to. I will instead argue that Grice's (non-)natural meaning distinction isn't exactly the right way to carve up the two senses, and that natural language semantics, without invoking a distinction about naturalness of meaning, already has the tools to differentiate between agentivity/episodicity in the meaning of the verb *to mean*. If these differences amount to an ambiguity, it's an ambiguity that is incredibly pervasive in English verbs and in verbs cross-linguistically. So, while much ado has been made about meaning and (non-)natural meaning in particular, there isn't anything particularly special about the alternations Grice observed for the English verb *to mean*.

In the next section, I'll talk about the morphosyntax of the English verb *to mean*, arguing that while it is not unusual in displaying argument alternations, the argument alternations it does display are relatively idiosyncratic. I'll then argue that the idiosyncracies of the argument alternations of *to mean* has been the cause of a great deal of confusion in the interpretation and defense of Grice's linguistic diagnostics. In §3 I turn to discussing the lexical semantic properties of *to mean*.

2 The (morpho-)syntax of meaning constructions

A well-known fact about (English²) verbs is that they divide into classes with respect to the argument alternations they allow (Levin, 1991). Two examples of argument-alternation classes are as follows:

- (10) a. The sun melted the ice. ERGATIVE VERBS
 b. The ice melted.
 (11) a. The picture frightened Dan. DYADIC UNACCUSATIVES
 b. Dan was frightened by the picture.

²For the next four sections, I will restrict my descriptive claims to claims about English. I will signal when this changes (in §6).

To melt exemplifies the class of ergative verbs; it can assign a theta-role to an agent subject and a theme object (as in (10-a)). But the verb *to melt* has another argument alternation, e.g. (10-b), in which the semantic theme can occur as the verb's surface subject. This is not true of non-ergative verbs like *met* (cf. *Magda met Matthew* and **Magda met*).

To frighten in (11) exemplifies a different argument-alternation class of verbs: dyadic unaccusatives are verbs whose agent subject (*the picture* in (11-a)) and theme object (*Dan* in (11-a)) can effectively switch places, linearly, with the theme becoming the surface object (just as it does with ergative verbs as in (10-b)). Unlike ergative verbs, the agentless form of dyadic unaccusatives can specify its agent, as in (11-b), but this agent can no longer be a syntactic argument of the verb; it must be introduced in a *by*-phrase (*by the picture*).

I believe it is a novel observation that *to mean* exemplifies an even quirkier class of verbs, of which *to cover* and *to amuse* are also members, in a way.³ Each of these verbs exhibit three argument alternations, which are exemplified below and schematized (for *to mean*) in (15).

- (12) a. The blanket covered the screen.
 b. Una covered the screen.
 c. Una covered the screen with the blanket.
- (13) a. Her antics amused Adam.
 b. Karen amused Adam.
 c. Karen amused Adam with her antics.
- (14) a. 'Knucklehead' meant Peter.
 b. Ernie meant Peter.
 c. By 'Knucklehead', Ernie meant Peter.
- (15) 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 μ that undergoes the action (but does not change state); and (iii) an instrument σ that is used to carry out the action. We know that the agent is an external argument of the verb, while the other two are internal arguments, in part because the triadic construction passivizes into an adjectival passive, not a verbal one (Grimshaw, 1990).

- (16) a. By 'Knucklehead', Peter was meant (by Ernie).
 b. *By 'Knucklehead', Peter is being meant (by Ernie).

The labels for the alternations in (15) are important. Unaccusative argument alternations (15-a) lack agents, even in some sort of deep-structure construal of the verb. In contrast, the agentive form in (15-b) has an agent, but merely lacks an instrument. Finally, the triadic form in (15-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 §5 I mention work on related phenomena in related classes of verbs. But I will note that this way of construing the argument alternations of *to mean* characterizes it as an underlyingly two-place predicate, with the agentive form exemplifying its true argument structure (i.e., with its agent x and theme μ being its external and internal argument, respectively). From this perspective, the instrument signal argument σ is a derived argument, introduced as a sort of add-on (in the triadic form in (15-c), by a preposition; Wilson 2020).

³Thanks to Ethan Poole (p.c.) and Beth Levin (p.c.) for help here; 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).

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 few 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 , syntactically it seems to select for a single external argument (the agent x) and a single internal argument (the signified meaning μ).

3 The (aspectual) semantics of meaning constructions

In this section, I discuss the lexical semantics of the verb *to mean*, and in particular how it interacts with tense/aspect to result in a range of possible interpretations (depending on various other things).

3.1 Verbs can be interpreted episodically or habitually, depending

What follows is a discussion of the episodic/habitual distinction that will become important for the reanalysis of Grice's linguistic tests that I present in §4. The basics of the distinction will be familiar to most, but it's important that we all get on the same page with respect to terminology.

I use the terms 'episodic' and 'habitual' to characterize a contrast in whether an eventuality is interpreted as occurring at least once (episodic) or in most cases (habitual). I take these terms to be synonymous with two other pairs of terms: eventive vs. stative, on the one hand, and stage-level vs. individual-level (Kratzer, 1995). If these terms diverge in the description of these data, it's not clear to me where.

Many verbs are in-principle ambiguous between episodic and habitual interpretations, e.g. *to work*.⁴ This semantic distinction is overtly marked in many languages, including African American English (Green, 2000), as exemplified in (17) and (18).

- | | | |
|------|--------------------------|-------------------------|
| (17) | a. Michael was working. | <i>episodic past</i> |
| | b. Michael been working. | <i>habitual past</i> |
| (18) | a. Michael is working. | <i>episodic present</i> |
| | b. Michael be working. | <i>habitual present</i> |

While Standard English doesn't overtly mark the aspectual distinction like African American English does, the distinction is nevertheless diagnosable in some cases in Standard English. The presence of certain eventive or temporal adverbs can force an episodic interpretation (19); and the Standard English simple present is almost always unambiguously habitual (in contrast to the present progressive, (20)).

- | | | |
|------|-----------------------------------|-------------------------------|
| (19) | a. Michael was working. | <i>ambiguous</i> |
| | b. Michael was working yesterday. | <i>unambiguously episodic</i> |
| (20) | a. Michael is working. | <i>ambiguous</i> |
| | b. Michael works. | <i>unambiguously habitual</i> |

The verb *to mean* is no different from *to work* in this respect; it can be made episodic with the inclusion of temporal adverbs, and is more naturally interpreted as habitual in the present tense.

- | | | |
|------|---|-------------------------------|
| (21) | a. By 'Knucklehead,' Ernie meant Peter. | <i>ambiguous</i> |
| | b. By 'Knucklehead,' Ernie meant Peter yesterday. | <i>unambiguously episodic</i> |
| (22) | a. By 'Knucklehead,' Ernie means Peter. | <i>habitual</i> |
| | b. ?By 'Knucklehead,' Ernie means Peter today. | <i>narrowly habitual</i> |

While (22-b) is acceptable, it is slightly awkward, as it requires forcing a habitual interpretation into a day-long span. In other words, (22-b) is acceptable in a context in which Ernie always refers to a different person each day with the term 'Knucklehead', and today he is habitually referring to Peter with that term.

⁴Some verbs, of course, are lexically stative, or more likely to receive an eventive interpretation than others. I discuss lexical stativity distinctions in communication verbs in §5.

3.2 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 instrument argument.

Recall that there are three semantic arguments of the verb *to mean*: an agent or speaker x ; an indicated meaning μ ; and an instrument used to signify, i.e. a signal, σ . From the point of view of semantic types, if a meaning relation has an agent, that agent is an individual, type e (an object over which an individual quantifier like *someone* or *something* can range). But it turns out that there is more than one type associated, in principle, with the instrument argument. In meaning relations with an agent (e.g. (21-a)), the signal is typically a linguistic string, or a sign, or a written word. But in meaning relations without an agent, the instrument or signal can be an eventuality.⁵

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).

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

As demonstrated in (23-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 DP subjects can denote eventualities, as illustrated in (23-b) and (23-c). It's important to note that, consistent with the Davidsonian tradition, these eventuality-denoting DPs are introducing eventualities in the semantics that are distinct from eventualities associated with the main verb. And they can be interpreted differently, too; in both (23-b) and (23-c), the eventuality introduced by the subject DP 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 a distinct aspectual properties.

We see this same property in meaning constructions. In particular, in meaning constructions without agents, the subject DP – the signal, or instrument argument – can denote an eventuality instead of an individual. And, in such cases, the eventuality can be interpreted episodically (as a single event) even while the main verb *means* is associated with a state, i.e. is interpreted habitually. This is illustrated in (24), altered slightly from (9).

- (24) a. Sam scratching his face means that he has a bad hand. *ambiguous*
b. Sam scratching his face meant that he had a bad hand. *non-natural only*

In (24-a), the subject DP – a gerund – is associated with an event. It has an episodic reading. In contrast, the main verb in (24-a) – in part because of its present-tense marking – is interpreted habitually, as a stative property. An accurate paraphrase of (24-a) is “The event of Sam scratching his face has the (stative) property of meaning that he has a bad hand.” To illustrate the variability we saw earlier, (24-b) has the same episodic, event-denoting subject as does (24-a), but contrastively (because of its past-tense marking), the eventuality associated with the main verb in (24-b) also has an episodic interpretation. (24-b) is most naturally accurately paraphrased 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. We need to attend to the lexical aspect of a given main verb (more on that in §5); we need to attend to the tense marking on that verb and to other components of the construction that favor a habitual or episodic interpretation of the main verb; and we need to, independently, consider the possibility that the subject denotes its own event. These all come into play in Grice's diagnostics, which I now turn to.

⁵In keeping with convention, I use the term ‘eventuality’ as a cover term for an event or a state (Moens and Steedman, 1988).

4 A new perspective on Grice's data

The core claims of this section are as follows:

- natural meaning is necessarily agentless and habitual
- non-natural meaning can have an agent, or not; and can be habitual, or not
- Grice's diagnostics for natural meaning are instead diagnostics for agentivity and episodicity, which come close to the (non-)natural meaning distinction, but don't perfectly capture it
 - Grice's **Passivizability** and **Agentivity** tests are tests for agentivity; and
 - Grice's **Actuality**, **Direct Quotability**, and **Factiness** tests are tests for episodicity

The differences between the (non-)natural meaning distinction, on the one hand, and the (non-)agentive and habitual/episodic differences, on the other hand, fall apart in certain contexts, and it is in those contexts that we can test the present claims.

4.1 Explaining agentivity differences in meaning constructions

I'll begin this subsection by focusing a bit more on my use of the term 'agent'. I intend to use the term in its syntactic sense, as the name of a particular thematic role, namely the role assigned to the cause or initiator of an event (for more discussion, see Dowty, 1991). While it is extremely tempting to associate this understanding of the term 'agent' with the philosophical or action-theoretic understanding of the same term, the two fall apart in certain cases. For instance, in sentences like *A met B*, the two individuals A and B arguably play an identical role in the meeting event, but are assigned different theta roles in the syntactic construction. That said, I will return at the end of this paper to discuss what can be made of the close-knit connection between the syntactic notion of agent and the philosophical one that underpins Grice's revolutionary pragmatic turn.

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.⁶

- (25) 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.

- (26) 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: these constructions necessarily marks an agent argument, and natural meaning does not have an agent. If fire means smoke 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. 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, imperfect test for non-natural meaning.

⁶In calculating these entailment relations, we need to hold fixed how we are interpreting them, i.e. we need to control for the episodic or habitual interpretation. To do this, it's best to insert a temporal modifier like *yesterday* into the sentence (to bring out the episodic interpretation).

To test this claim, we cannot look for an instance of natural meaning with an agent, but we can find an instance of non-natural meaning without an agent, as in (27).

(27) The smoke alarm means fire. → Somebody or other meant fire by the smoke alarm.

By virtue of the fact that it lacks an overt or specified agent (given its status as an unaccusative form), (27) does not pass the Agentivity test. This is, of course, despite it encoding a non-natural meaning relation.

Grice's Passivizability test, too, is a test for agentivity. 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 (28-a), and a passivized version of a triadic form (28-b).

(28) a. 'Knucklehead' means Peter. *unaccusative form*
b. Peter is 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 or instrumental 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 (54-b), in contrast, is a passivized form of the triadic configuration *By 'Knucklehead,' someone meant Peter*, i.e. ' μ was meant by σ '.⁷

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 has an implicit agent. This is a familiar claim from the transitive sentence in (29-a) and its passivized version in (29-b).

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

The source of the controversy surrounding Nixon's utterance of (29-b) is that it entails someone made the mistakes, by virtue of its relation to the transitive sentence in (29-a). But this property of passive constructions is not a property of unaccusative constructions like (28-a).⁸

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. (30) is an example of a natural meaning construction failing the passivizability test (as Grice predicts), but (31) is an example of a non-natural meaning construction also failing the passivizability test, in contrast to Grice's predictions.

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

The unaccusative form in (31-a) cannot be passivized because it, too, lacks an agent.

⁷By virtue of the fact that it specifies an agent, the agentive form, too, can be passivized: *Ernie meant Peter* can be passivized to the somewhat awkward *Peter was meant*. So while passivizability isn't a sole property of agentive forms, I will focus here on the triadic form.

⁸There are a variety of ways that this implicit agent argument is said to be able to be interpreted in context (Cinque 1988, Collins *to appear*): definite/anaphoric (to a determinate, identifiable agent); existential (to a determinate, non-identifiable agent); and generic (to everyone in the context of evaluation). However, it's not clear to me that there is a clear, robust distinction between the definite/anaphoric and existential uses.

4.2 Explaining episodicity 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: episodicity. The correlation between non-natural meaning and episodicity, 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 or instrument argument σ is eventive. This is slightly different than a claim about the habitualness or episodicity of the meaning relation, but it's correlated. In particular: for natural meaning, the meaning relation is always habitual, but the signal or instrument argument of this habitual relation can be eventive. For non-natural meaning, the meaning relation can be habitual or episodic, and when it is episodic, its signal must be eventive. I'll explain in detail in what follows.

Recall from §3.2 that the signal or instrument argument σ of meaning relations can denote an eventuality instead of an individual. This is illustrated in (32), repeated from (9).

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

In (32-a), the instrument subject, *Sam scratching his face*, is event-denoting; in (32-b), a triadic form, the instrument is encoded in a *by*-phrase, and the subject is now an individual (the agent of the meaning relation).

4.2.1 Direct Quotativity is a test for event signals

Insofar as we have judgments about the acceptability of quotation marks, Grice's **Direct Quotability** test is a diagnostic for whether or not a signal or instrument argument σ in a meaning relation is event-denoting (as opposed to individual- or state-denoting). In particular, we can report a meaning μ in quotation marks if 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 (33) (from (7)).

- (33) 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 or instrument 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 (34) 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 (35) are instances of non-natural meaning, and (contrary to Grice's generalization) are unacceptable with quotation μ arguments.

- (34) **natural meaning** constructions with eventive signals
a. Smoke (billowing) means, 'There's a fire!'
b. Flowers blooming means, 'It's spring!'
(35) **non-natural meaning** constructions with non-eventive signals
a. *The tricolor flag means, 'Croatia!'
b. *A pie in the window means, 'The spy ring has been compromised!'

In (34), the natural meaning signals are made eventive by virtue of the fact that they're encoded in a gerundive DP. As a result, the μ or theme argument of these natural-meaning relations can appear in quotes. In contrast, in (35), the DP 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 eventive, 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 (34) and (35) 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.

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

Crucially, 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).⁹ There are at least two ways to condition an actuality implication for an event-denoting argument: 1) include proximal marking in the event-denoting DP; 2) place the event-denoting DP in a construction that is interpreted episodically. I will illustrate each in turn.

The sentences in (36) encode natural meaning relations. In (36-a) the signal σ is an event-denoting DP, the event of lightning striking. In the case of (36-a), because the verb *to mean* is in the present tense, this meaning relation is most naturally interpreted as habitual (see §3.1).

- (36) a. Lightning striking means a storm is near.
 b. That lightning striking means a storm is near.
 c. 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 (36-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 habitually: the meaning relation is still interpreted as non-contingent, holding over time. It's just that (36-b) expresses that the very recent and actual lightning strike is an instantiation of that habitual meaning, especially in relation to the meaning, which itself can be interpreted either generically or with respect to the here-and-now. 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 (36-c), we achieve an actuality implication on the event-denoting argument in a different way: by making the main verb episodic, instead of habitual. This was accomplished, as foreshadowed in §3.1, by changing the present-tense marking on the verb to past tense. (Again, this change doesn't *require* an episodic interpretation, but it favors it.) If we are instead looking at a meaning relation that is episodic, 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 that its occurrence was associated with the meaning μ the speaker attributes to it.

So, 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 episodic. 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 episodically. If a meaning relation is interpreted episodically, it is necessarily a non-natural meaning relation (i.e. it is contingent).

Two additional notes here: one, this is an asymmetrical entailment, i.e. non-natural meaning relations don't *need* to be episodic, and thereby don't need to have actual eventive arguments. And two, there's a way to

⁹It is more typical to refer to this as an actuality entailment, not an actuality implicature. 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.

achieve an actuality implication that doesn't involve an episodic 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 (37).

- (37) 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*
- (38) Ring-shaped spots mean measles, but no one here has {ring-shaped spots/measles}.

Contrast (38) with Grice's original test in (37-a): the instrument argument in (37-a) contains proximal marking, which prompts an actuality implication for the event associated with the DP, resulting in Grice's original intuition that the sentence entails that the person in question has spots (and therefore, given the habitualness of the meaning relation, that he has measles). In contrast the instrument argument in (38) 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 (39), we can test how event actuality affects 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 (39-a) is a parallel of the Grice example in (37-b): it's a non-natural meaning construction that passes the non-actuality test, i.e. that lacks an actuality implication.

- (39) 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 (39-a) is the tense on the verb: it is in present tense, which means it's most naturally interpreted habitually. But when we change the tense to past tense (cf. also (9-b)), we switch to an episodic 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.

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 (40), 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.

- (40) 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 agentivity and episodicity 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 necessarily habitual, 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. This is schematized in Table 2.

	<i>description</i>	agentive?	~episodic eventive?	actual?
Passivizability (a)	Can ' σ means μ ' be passivized?	✓	n/a	n/a
Agentivity (b)	Does ' σ means μ ' entail 'Somebody means μ '?	✓	n/a	n/a
Actuality (c)	Is the speaker committed to μ being actual?	n/a	✓	n/a
Direct Quotability (d)	Can the meaning μ argument occur in quotation marks?	n/a	✓	✓
Factiness (e)	Can the signal σ be paraphrased with 'the fact that'?	n/a	✓	✓

Table 2: What Grice's diagnostics really test for

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 one directly correlated property (agentivity), and another loosely correlated property (eventivity). In the rest of this paper, I address the question of how broad or general this phenomenon is. In particular, in §5 I argue that this agentive and episodic/eventive distinction varies in an interesting way within English across the class of meaning verbs, and has already been documented in studies of verbs of communication more broadly in English. In §6, I argue that, while many languages are like English in having a single lexical item *to mean*, other languages have two distinct lexical entries, and that this lexical distinction tracks agentivity and episodicity, rather than (non-)natural meaning.

5 The role of agentivity and episodicity in other English verbs

There are several English verbs that pattern like *to mean* in exhibiting the argument alternation pattern in (15) that is unique to the *cover/amuse* class of verbs. They are, as far as I can tell: *to denote*; *to designate*; *to indicate*; *to mean*; and *to refer to*.

- (41) 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 semantically independent. There are some synonyms of *to mean* that take on a different meaning in the absence of the signal argument, as demonstrated in (42).

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

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

- (43) *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:

- (44) **Agentivity**
 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*
- (45) **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 habitual or episodic, de-

pending 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 episodic relation, and can, in contrast, head an imperative (e.g. *Refer to Peter (when you use the word 'Knucklehead')!*).

Given this variation in lexical aspect, we would expect Grice's tests for eventivity or episodicity 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 episodic interpretations across the board, regardless of the (non-)naturalness of the meaning relation. And this is what we see:

- (46) 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 habitual interpretations, and therefore any sort of natural meaning relation. And this is what we see:

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

I'll also add that there is now a small cottage industry in the typologizing of sentence-embedding English verbs, and these studies indicate that the habitual/episodic semantic alternations that were pivotal above are quite common. Anand et al. (2017) argue that a group of terms they dub 'communicative predicates' behave differently from other sentence-embedding verbs (doxastic, emotive, inferential) in that they allow for certain inanimate objects in the instrument role, and in that they are never factive (on the μ argument). In a series of papers (Major 2021, Major *to appear*, Major and Stockwell *to appear*), Travis Major focuses on some unique syntactic properties of the verb *to say* and its correlates in other languages, observing that *to say* is systematically ambiguous between eventive and stative interpretations, and that this alternation has syntactic reflexes for e.g. argument structure and aspectual marking, much like the reflexes reviewed above.

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 instrument argument); and its aspectual polysemy (in particular, its ability to be interpreted habitually or episodically, depending on the nature of its arguments, and the morphosyntactic properties of the construction).

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 and, I believe, still a commendable and welcome contribution.

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 episodicity, not (non-)natural meaning.

6 The role of agentivity and episodicity 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 overwhelmingly many 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.

I'll just caution here that things get a little murky, because it's quite common for languages to employ multiple synonyms or near-synonyms for *to mean*, just as English does (cf. §5). So simply asking a speaker of a language how many words they have for *to mean* could reveal a variety of different, slightly irrelevant lexical differences, including differences in argument structure or lexical aspect. To make the claims I'm making here, I tracked down languages for which one word could be used in a canonical natural meaning construction, but not in a canonical non-natural meaning construction; and another word that could be used in a non-natural construction but not a natural one. The goal of this section is to argue that, for these languages, the lexical distinction actually tracks agentivity (and, to a lesser extent, episodicity).

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

(48) **Greek**

- 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.'

The verb in (48-a) is what is elicited for natural meaning relations, and the verb in (48-b) is what is elicited for non-natural meaning relations. 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 (49).

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

This construction uses the unaccusative, non-agentive verb *simeni*, rather than the agentive verb *enoi*. Speakers of several other languages report the same configuration, including Estonian, Icelandic, and Romanian.

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.¹⁰

(50) **Spanish**

- 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.'

It also has a second phrase, 'wants to say,' that requires a syntactic agent. (51) exemplifies the Spanish version of the English triadic construction.

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

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.

¹⁰It is possible to associate the verb *significar* with an experiencer, as in *Para mí* ('For me')... But this is different than an agent argument, syntactically and in other ways. I'll return to discuss experiencer arguments in §7.

- (52) Fuego quiere decir corre.
 fire want-3SG to-say run
 ‘Fire means run!’ (lit. ‘Fire wants to say run.’)

The use of the volitional *querer decir* in (51) is only appropriate in a subset of conditions that (50-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’ construction 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).¹¹

- (53) **Avatime**
 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.”

Both versions of ‘to mean’ in Uyghur (a Turkic language spoken in China) involve the verb ‘to say,’ but the language distinguishes between episodic and habitual meaning relations via other means: an episodic (non-natural) meaning relation in (54-a) involves a ‘meaning’ nominal, amounting to something like ‘The word’s meaning when one says ‘bodek’ is fat’.¹²

- (54) **Uyghur**
 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.’)

While habitual meaning relations (e.g. (54-b)) involve a subjectless use of ‘say’ as an impersonal marker that is restricted to habitual or generic contexts (see Major *to appear* for more detail on this construction).

In sum, lots of languages employ a variety of different words or strategies for specifying meaning relations. When a language has more than one word for ‘to mean,’ and when this distinction comes and goes with the (non-)natural meaning distinction, I’ve provided preliminary evidence here that the distinction is really tracking something like agentivity (in the case of Spanish) or episodicity/habitualness (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 episodically or habitually. In the concluding section that follows, I speculate about the relationship between these linguistic properties and the more metaphysical notions Grice advocated for.

7 Conclusions and extensions

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

¹¹Thanks to Divine Agyepong and Vincent Azafokpe via Travis Major for the Avatime data.

¹²Thanks to Gulnar Eziz and Akbar Amat via Travis Major for the Uyghur data.

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.

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 episodically, or contingently (what I've called 'episodicity'). If we take the reasonable position that natural meaning, just intrinsically, cannot have be specified with an agentive argument, and cannot be interpreted episodically, I argue that Grice's linguistic observations are accounted for.

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 verb *to mean* as ambiguous or polysemous or having more than one sense, at least in contrast to any other verb of communication or signification. These verbs – across English and cross-linguistically – can be interpreted episodically 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 episodicity, 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 tentatively concluded that there are two and only two lexicalized syntactic arguments of *to mean*, the agent x and the meaning μ . But two other syntactic arguments can be associated with the meaning relation: the signal σ , introduced in the accusative form (55-a) or by the preposition *by* (55-b), or an experiencer, introduced by the preposition *to* (55-c).¹³

- (55) 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 meaning μ ; ii) an agent x ; iii) a signal σ ; and iv) an experiencer y . This is schematized in (56).¹⁴

- (56) 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.

¹³Crucially, 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.

¹⁴The instrument 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 (55-a). This is not true of the experiencer, so we consider that a non-core argument, or an adjunct.

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 (57)-(59), I do this for implicit agent arguments.

- (57) **common terms**
 A: *blah blah blah* word.
 B: #What does ‘word’ mean?
 B’: #What do you mean by ‘word’?
- (58) **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’?
- (59) **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 *word*, there is no variation in understanding from the general population to any sort of specialist population. It is therefore weird to ask what ‘word’ means (cf. (57)B), but it is also weird to ask what someone means by ‘word’ (as in (57)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 felicitously ask about the meaning of the term writ large.

In (60)-(62), 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.

- (60) **common terms**
 A: ‘Word’ means X.
 B: No, it means Y!...
 B: ...You’re wrong.
- (61) **specialist terms** (where *B* is a non-specialist)
 A: ‘Morpheme’ means X.
 B: #No, it means Y!...
 B: #...You’re wrong.
- (62) **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.

Conflict of interest

Conflict of interest: No

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.
- Cinque, G. (1988). On *si* constructions and the theory of ARB. *Linguistic Inquiry*, 19:521–581.
- Collins, C. (to appear). *Principles of argument structure: a Merge-based approach*. MIT Press.
- Comrie, B. (1976). The syntax of action nominals. *Lingua*, 40:177–201.
- Dowty, D. (1991). Thematic proto-roles and argument selection. *Language*, 67:547–619.
- Ginzburg, J. and Sag, I. (2001). *Interrogative investigations: the form, meaning, and use of English interrogatives*. CSLI Stanford.
- Green, L. (2000). Aspectual *be*-type constructions and coercion in African American English. *Natural Language Semantics*, 8:1–25.
- 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. (1990). *Argument structure*. MIT Press.
- Kratzer, A. (1995). Stage-level and individual-level predicates. In Carlson, G. and Pelletier, F., editors, *The generic book*. Chicago.
- Levin, B. (1991). *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. (2021). On the nature of 'say' complementation.
- Major, T. (to appear). Re-analyzing 'say'-complementation: implications for case theory and beyond. *Natural Language and Linguistic Theory*.
- Major, T. and Stockwell, R. (to appear). "say"-ing without a voice. In *Proceedings of the North East Linguistics Society (NELS) 51*. UMass GLSA.
- 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.
- 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.
- Wilson, M. (2020). The reversible core of object experiencer, location, and *govern*-type verbs. In *Proceedings of NELS*, volume 49, pages 285–294.