

All Update is Goat Update
Andy Egan and Timothy Sundell¹
Draft of June 18, 2024

1. Introduction: The ridiculous informational richness of speech acts

Five minutes before the appointed starting time of Carrie’s talk at the Central APA, Alphie says to Betty, in a loud, clear voice in a midwestern-american accent,

(A) “Professor, she’ll be here 10 minutes from now”.

There is a strikingly rich array of messages that Alphie could, given the right background conditions, communicate to Betty by means of this utterance. For a few examples among many, consider (1)-(26).²

1. Carrie will be in the ballroom of the Palmer House Hilton at 1:05pm.
2. Alphie believes that Carrie will be in the ballroom of the Palmer House Hilton at 1:05pm.
3. Alphie has compelling evidence that Carrie will be in the ballroom of the Palmer House Hilton at 1:05pm.
4. Whether Carrie will be in the ballroom of the Palmer House Hilton at 1:05pm is relevant for the purposes of Alphie and Betty’s conversation.
5. Alphie believes that Betty didn’t already know that Carrie will be in the ballroom of the Palmer House Hilton at 1:05pm.
6. Carrie is not presently in the ballroom of the Palmer House Hilton.
7. Carrie will be able to give her talk.
8. Carrie will be late for her talk.
9. Carrie is not in San Francisco.
10. Betty should stall, rather than canceling the session and asking everybody to leave the ballroom.
11. Alphie believes that whether Carrie will be in the ballroom of the Palmer House Hilton at 1:05pm is relevant for the purposes of Alphie and Betty’s conversation.

¹ Thanks to SO MANY PEOPLE. Chris Barker, Elisabeth Camp, Dan Harris, Karen Lewis, Craige Roberts, Mandy Simons, Michael Glanzberg, Dan Lassiter, Mike Titelbaum, John Perry, Wolfgang Schwartz, Stacie Friend, Matthew Chrisman, Henry Schiller, Jessica Rett, Shannon Bain, Paul Pietroski, Eliot Michaelson, Luke Elmore, participants in MIT/NYU philosophy/linguistics graduate seminar Spring 2022, participants in John Perry’s Zoom Group in April 2024, audience at colloquium talk at University of Edinburgh...

² This example is a variation on one from (Sundell, 2009).

12. Whoever Alpie is referring to with “she” will be in the ballroom of the Palmer House Hilton at 1:05.
13. Carrie will be in the Palmer House Hilton ballroom ten minutes after the time of Alpie’s utterance.
14. Alpie believes that Carrie will be in the Palmer House Hilton ballroom ten minutes after the time of Alpie’s utterance.
15. The fact that Carrie will be in the Palmer House Hilton ballroom ten minutes after the time of Alpie’s utterance is relevant to some goal or project that Alpie and Betty are pursuing.
16. Alpie believes that the fact that Carrie will be in the Palmer House Hilton ballroom ten minutes after the time of Alpie’s utterance is relevant to some goal or project that Alpie and Betty are pursuing.
17. There are a proposition P, person x, and location l such that: (i) P is the proposition that x will be at l at 1:05pm; (ii) x is the person referred to by Alpie’s use of ‘she’; (iii) l is the location referred to by Alpie’s use of “here”; (iv) Alpie’s utterance expressed P, and (v) P is true.
18. There are a proposition P, person x, time t, and location l such that: (i) P is the proposition that x will be at l ten minutes after t; (ii) x is the person referred to by Alpie’s use of ‘she’; (iii) l is the location referred to by Alpie’s use of ‘here’; (iv) t is the time referred to by Alpie’s use of ‘now’; (v) Alpie’s utterance expressed P, and (vi) P is true.
19. There is some proposition P such that (i) P is the proposition expressed by Alpie’s utterance, and (ii) P is true.
20. Alpie is in a position of lower social standing or authority than Betty.
21. Alpie takes himself to be in a position of lower social standing or authority than Betty.
22. Alpie is concerned to explicitly signal that he takes himself to be in a position of lower social standing or authority than Betty.
23. Alpie speaks English.
24. Alpie believes that Betty speaks English.
25. Alpie is not from Australia.
26. Alpie does not have laryngitis.

The potential messages here range from information based almost entirely on the literal, linguistically encoded content of Alpie’s utterance to information that has nothing at all to do with that content, with a great deal of information falling somewhere in between. Likewise, the messages themselves vary in content from propositions that concern, exclusively, matters to do

with Carrie, her location, and her talk, to propositions that concern, exclusively, Alpie, Betty, and their immediate conversational context, again with a great deal in between.

These are pretty clearly all pieces of information that Betty could extract from Alpie's utterance, and update her beliefs with. They are also all pieces of information that Alpie could predict Betty will, and intend Betty to, extract and update her beliefs with. And they are all pieces of information that could, as a result of Alpie's utterance, become part of the common ground of Alpie's and Betty's conversation.

A natural thought to have is that this wide diversity of communicative effects must be the product of a wide diversity of communicative mechanisms. That, for example, the mechanism underwriting communication of the literal content of the sentence Alpie uttered in the context in which he uttered it (proposition 1, or something very like it) is very different from the mechanism underwriting communication of (for example) propositions 2, 4-8, 12-19, 20-22, or 26, and that the mechanisms underwriting communication of propositions 2, 4-8, 12-18, 20-22, and 26 are themselves quite different from each other.

For example, it's natural to divide them up this way: Betty's update with (1) is the upshot of the standard mechanism of literal assertion and acceptance, and is (assuming other factors cooperate) an instance of the transmission of knowledge by testimony. Her update with (2) is the result of an importantly different process of pragmatic inference. Updates with (4)-(8) are the results of the distinctive kind(s) of pragmatic inference that generate(s) conversational implicatures. An update with (19) would be the paradigmatic result of Stalnakerian (1978) diagonal repair. Updates with (17) or (18) would be cases of updating with the sort of reflexive truth conditions Perry and co-authors talk about.³ Updating with (20-22) will be the result of some distinctive mechanism associated with honorifics, and perhaps also pragmatic inferences that are parasitic on those mechanisms. And the update with (26) is a "goat update" - an update with a proposition about manifest features of observable events, like how, if we see a goat walk into the room, we update with the fact that a goat walked into the room. (Or with the fact that the goat that just walked into the room is large, or hairy, or smelly, or noisy.)⁴ And literal assertion, implicature, diagonalization, updating with reflexive truth-conditions, and goat update are all, of course, distinct mechanisms.

³ See for example (De Ponte et al., 2023; Korta & Perry, 2011; Perry, 1988, 2001).

⁴ The term "goat update" is due to a justly famous example from (Stalnaker, 1978).

Goat updates in particular are standardly introduced precisely in order to distinguish them as different in kind from more distinctively language-based sorts of updates—there's on the one hand the sort of update that happens when we accept the content of a speaker's assertion, and on the other hand the importantly different sort of update that happens when we all notice each other noticing a goat ostentatiously walking in to the room.

We are going to make the case for a deflationary view, according to which all of these quite diverse-looking communicative effects are the result of a single, unified mechanism. All of the rich array of communicated messages and updates associated with an utterance are the results of updating with the manifest facts of and about the utterance, while holding fixed various background assumptions.

Our slogan: All update is goat update. A goat walks in and we all update with the proposition that a goat walked in. Of course, we don't *just* update with the proposition that a goat walked in. Depending on what else we believe, we are likely to form a number of other beliefs as well, which are consequences of the proposition that a goat walked in, together with various of our other beliefs. For example, depending on what else we believe, we might also come to believe that the surrounding environment is at least minimally hospitable to goats. Or that there is a goatherd somewhere nearby. Or that the nearby goatherd is not maximally attentive. And so on.

Moreover, we may well take note of each other taking note of the fact that the goat walked in, and so all take each other to have updated with the proposition that a goat walked in. And we may also take each other to take each other to have so updated, etc. In that case, *that a goat walked in* becomes part of the conversational common ground. (This is what's happening in the original goat example in (Stalnaker, 1978).) Lots of other things are likely to enter the common ground as well, depending on what assumptions, if any, we are collectively holding fixed about the features any goat that walked in the room would have. For example: That one of Farmer Bob's goats walked in. That goats are allowed in the seminar room. That something has gone wrong with Farmer Bob's goat-containment systems. That the goats' food trough is empty. That Farmer Bob wants us to leave the seminar room. And so on, depending on the precise nature of the common ground in the conversation at the time of the goat's conspicuous entrance.

Utterances are, we propose, no different. There's the event of the utterance. And there are a lot of assumptions we might make (either individually or collectively) about features of the

utterance. Assumptions about truth value, about lexical meaning, about the resolution of context dependence, about relevance, about the appropriateness of the quantity of information offered to the conversational purposes at hand, about the speaker's sincerity, about the utterance's reflectiveness of the speaker's native accent or social status, or of the condition of their vocal apparatus, etc., etc....

The central thing that happens in our example is that Alphonse makes an utterance, which has some observable features, and Betty updates with the fact that Alphonse made an utterance with those features. Then, depending on what else Betty believes about Alphonse, his utterance, and the surrounding context, she is likely to be able to extract a great deal of additional information from the manifest fact that the utterance occurred, and had the observable features that it had.

Based on the acoustic features of Alphonse's utterance, if Betty knows the relevant facts about the pronunciation of words in English and English grammar, she can get *that Alphonse produced sentence S*. If she knows, in addition, the relevant facts about English word meanings and the facts about context relevant to resolving any context-dependence in S, she can get *that Alphonse expressed P*, where P is the proposition semantically expressed by S in Alphonse's context. If she assumes that Alphonse is speaking truly, then she can get to P itself. If she believes further that Alphonse is speaking cooperatively along various dimensions, then she can get to a range of implicatures of Alphonse's having said that P. And, depending on what else she believes, she is likely to be able to extract a great deal more information than this.⁵

Terminological note: We'll use the terms "goatism" and "informationalism" to describe the sort of picture we're advocating for, because we think that they highlight different features of the view that it's useful, in different contexts, to be reminded of. (And also, frankly, because it's occasionally nice to have a more respectable name for the view than "goatism".)

A theoretical benefit of framing an account of communication in these goatism, informational terms is that doing so gives us a unified story to tell about the whole rich array

⁵ Above, and also repeatedly below, we will move not-super-carefully between talking about knowing, assuming, and believing various things about speakers, utterances, and circumstances, and using one's knowledge/assumptions/beliefs to extract additional information from the fact of (and facts about) an utterance. In all cases the central load bearing notion for us is going to be *belief*. Sometimes it's convenient to talk about knowledge, because the beliefs that are doing the inference-supporting work are ones that are known. Sometimes it's convenient to talk about assumptions, because the beliefs doing the inference-supporting work are ones that are taken for granted in an assumption-y way. We will assume, without argument and perhaps contentiously, that knowing and assuming are ways of believing. If that turns out to be too contentious, we will cede the terminological ground and do a search-and-replace so that we are talking everywhere about belief.

of communicative effects illustrated by (1)-(26). They're all based on extraction of information from the fact of (and facts about) the utterance, by drawing inferences from those facts together with some background assumptions. And speakers can set hearers up to engage in that process of information-extraction in ways that are, at least potentially, completely predictable, intentional, and transparent.

A rich array of potential background assumptions gives rise to the rich array of potential communicative effects with which we began. And a correspondingly rich array of assumptions for speakers to make about the assumptions hearers are making, and of intentions speakers might have about exploiting those assumptions for communicative effect, gives rise to a rich array of communicative intentions that a speaker might have in producing an utterance. And any of these intentions might be recognized by hearers, and might be predicted and intended by speakers to be recognized by hearers.

There is, in this respect, nothing special about linguistic utterances. There is a similarly rich array of assumptions that goatherds might make about which assumptions witnesses might make about the nature and origin of any goats walking into the room, which in the right circumstances could also be exploited for a wide range of communicative purposes. And obviously there is nothing special about goats. The same is true of any observable event.

The resulting picture is one on which, viewed from a certain altitude, there is nothing very distinctive about the way we learn things from other people's words. The ridiculous informational richness of speech acts is a special case of the ridiculous informational richness of everything. We learn from speech acts in the same way we learn from anything else. We update with the fact of, and facts about, some manifest event. And we take advantage of other information that we have at our disposal to leverage additional information from the manifest event. A speaker with the right kinds of knowledge about the information a hearer will be bringing to bear can use that knowledge to leverage additional communicative force into their utterance. But that's not a distinctive feature of utterances—an interlocutor with the right kind of knowledge about the information we'll bring to bear in thinking about their actions generally can use that knowledge to leverage additional communicative force into their winks, nudges, arm-wavings, head shakings, up-and-down-jumpings, pratfalls, laughs, grunts, foot tappings, eyebrow wiggles, goat-releasings, and whatever other actions you like.

The ‘viewed from a certain altitude’ qualification above is important, though. In another respect, this picture is compatible with various aspects of linguistic communication being special, distinctive, and in need of distinctive branches of theorizing. The sort of information that Tim brings to bear in extracting (for example) the message *that it’s snowing in Baltimore* from Andy’s utterance of “it’s snowing in Baltimore”—information about the ways in which English words are pronounced, how they combine into sentences, and how the meanings of the words compose to form sentence meanings—is distinctive, and the story about the cognitive mechanisms by means of which Tim brings that information to bear is likely to be distinctive as well. The goatist view is *not* intended as a proposal about the specific architecture of the cognitive mechanisms underwriting linguistic competence, and its correctness is compatible with a wide range of possibilities about the mechanisms by which the relevant information is encoded and brought to bear. Rather, the view is intended to bring out the structural unity among the many different kinds of inferences we draw on the basis of an utterance, however it turns out the relevant beliefs are encoded or how the relevant inferences are carried out.

The remainder of this paper is structured to respond to the three most common sorts of negative reactions we have received when presenting this view to people: “yeah, duh”, “no way”, and “so what?”. We preview below, and then elaborate in sections 2, 3, and 4.

“Yeah, duh”:

Of course that’s how it works, at a high enough level of abstraction. Obviously that’s how belief update *always* happens—we update with the fact that something happened, and then we draw additional inferences based on that fact plus other stuff we believe. (In a Bayesian framework: We update by going to high credence on some evidence (1, if we’re conditionalizing, lower than that if we’re Jeffrey conditionalizing) and then, depending on what our priors are like—depending, that is, on how our credence was antecedently distributed over various possible ways for things to be—that update brings us to high credence on a number of other things we may not have been high credence on before.) How else could it possibly work, how could anybody ever have been committed to thinking it was otherwise, and how could it be of any theoretical interest to make a big fuss about this obvious and undeniable fact?

Our response here will have two components: First, presenting an alternative picture of how communication might work which is incompatible with goatism (thus demonstrating that goatism is not the only possible picture). Second, making the case that explicitly framing things in goatist terms can still be theoretically interesting and helpful, even for those who have always been at least tacitly committed to its correctness.⁶ We make a general version of this second point in the “yeah, duh” section, and then—in the “so what?” section—work through several specific linguistic phenomena in terms of the flat, goatist picture in detail. We think some interesting and helpful lessons emerge.

“No way!”

That’s not enough, and/or not the right kind of, theoretical resources to explain all of the things we need to explain. For example, this kind of picture trades only in belief, and the only kinds of update it’s going to be able to model are doxastic updates. But there are lots of non-doxastic updates that it’s important to be able to capture in our theories of communication and of linguistic meaning. It’s very plausible, for example, that the updates performed in response to imperatives and questions are non-doxastic, or at least not entirely doxastic (REFs). Similarly for the updates that feature in dynamic semantic accounts—introduction of discourse referents (REFs), tests on the context set ((Groeneveld & Veltman, 1997; Veltman, 1996; Willer, 2013)), filters of the context set that aren’t intersective updates ((Barker, 2002, 2009, 2012)), updates to speakers’ discourse commitments (REFs), etc. These kinds of updates go beyond doxastic updates in at least two different ways. For one thing, they are in the first instance characterized as updates to *contexts*, not to states of individual hearers. For another, even if the relevant features of context are reducible somehow or other to psychological features of the individual participants in the conversation, it’s not at all obvious—and not really very plausible—that they’ll reduce entirely to *doxastic* features of participants in the conversation.

⁶ Autobiography: We think we were tacitly committed to this kind of picture all along, and probably never committed ourselves to anything incompatible with it. But we definitely were not always explicitly framing things in these terms in our thinking or writing. We’ve found it clarifying and fruitful in many cases, with respect to many issues. Even if our endorsement of the flat, “it’s all goat update” picture isn’t a new belief (as good Stalnakerians we’re committed to thinking it’s not, since it’s entailed by things we already believed), it’s useful to draw out and make explicit so we can use it as a framework for thinking about particular issues, questions and phenomena.

The response here requires (a) being a bit more careful about how we state our thesis, and (b) introducing a type of move that will be central to the goatist program, explaining how goatism-friendly doxastic updates based on the observed fact of and observed facts about an utterance, together with knowledge of the conventional communicative roles of sentences (more generally, expressions) of a language, can systematically bring non-doxastic updates in their wake.⁷ We provide details in section 3.

“So what?”

Above, we emphasized the qualifier, “from a certain altitude”. But from a high enough altitude, everything starts to look the same. Supposing the goatist thesis is true, does it actually shed light on phenomena of theoretical interest? We think it does. The most obvious selling point for the informationalist framework is unification. An informationalist approach promises to offer a unified account of the enormous, initially quite diverse-looking, range of communicative effects of a very broad, initially quite diverse-looking, range of communicative acts. And it brings the whole range of communicative effects under a common theoretical umbrella. Not trivially, by needlessly abstracting away from details, but in a way that promises to clarify how the various effects are related to each other. This in turn makes it easier to see, and easier to theorize about, some phenomena where several different kinds of communicative effects interact with each other.

We’ll talk about a few different phenomena to substantiate this. But our headline example will be putting the goatist framework to use by applying it to a class of phenomena in which the cognitive or communicative significance of an utterance comes apart from the proposition the utterance semantically expresses in its context, often—though not always—because the utterance is only imperfectly understood by its intended audience, or because part of what’s at issue in the conversation is the meanings of some of the words being used. This is a class of phenomena that

⁷ We don’t, officially, place any bets on the communicative roles of expressions of a language being grounded in convention. (Nor on the correctness of any particular account of convention.) It’s important to us (one of us more emphatically than the other) that the account not rely in any ineliminable way on meaning being grounded in community-level, external factors like the conventions in place in a large-scale speech community or public language. While many of our examples will involve reference to conventional meanings and public languages, we take it they could easily be reframed in terms of speakers’ beliefs about one another’s individually-determined meanings, and in fact we’ll help ourselves to more individualist or internalist modes of description in examples that are more naturally presented in that way.

have already been quite broadly theorized about under different modes of presentation, probably most prominently by John Perry and Robert Stalnaker⁸, and more recently by (Barker, 2002, 2012; Plunkett & Sundell, 2013; Rudolph, 2021; Rudolph et al., 2020; Rudolph & Kocurek, 2021; Sundell, 2014, 2016). These, according to us, are all instances of a single overarching phenomenon, and framing our discussion of it in explicitly goatist terms is helpful for making clear what the overarching phenomenon is and how it works, why and in what respects each of these theoretical modes of presentation are enlightening, and why and in what respects they're potentially misleading. We think the resulting picture is helpful and clarifying, and we hope you will too.

A bit of conceptual ground clearing and terminological stipulation before we move on: In what follows, we'll talk a lot about the information that audiences bring to bear (and that speakers expect them to bring to bear) when updating in response to an utterance. We'll use a number of terms to talk about this—for example “assumptions”, “background knowledge”, “beliefs”, “information”. In some cases, the belief in question counts intuitively as knowledge, and so when it won't lead to confusion, we'll describe it that way, but the crucial load-bearing phenomenon is always going to be the belief, not whatever additional features the belief has that lead it to qualify as knowledge.⁹

It's also important to emphasize that we are adopting a very permissive Stalnakerian notion of *belief*, according to which there is a lot of tacit belief, and potentially a lot of diversity in the sorts of particular cognitive structures underwriting different beliefs. What's important to counting as believing P, for our purposes, is being in the right functional/dispositional state, where those states are characterized at a quite high level of abstraction, compatible (by design) with a great deal of multiple realizability in implementation within and across believers.¹⁰

⁸See for example (De Ponte et al., 2023; Korta & Perry, 2011; Perry, 1988, 2001; Stalnaker, 1978, 1987, 2006, 2009)

⁹ We put it this way because, as old-fashioned knowledge-is-belief-plus-something-else philosophers, it's the natural way for us to say it. But we don't want to pick a fight with Williamsonian (Williamson, 2000) knowledge-is-prime philosophers if we can help it. So we're happy to retreat to something like, “the load-bearing phenomenon is belief, not the more-demanding-even-though-not-factorizable-into-belief-plus state of knowledge”, or whatever the primacy-of-knowledge-compatible restatement of our hedge here turns out to be. (Though see (Jackson, 2002) for an expression of a sort of skepticism about non-factorizability that we are sympathetic to.)

¹⁰ So that beliefs are the states that, for example, indicate that P, dispose their bearers, to act in ways that would satisfy one's desires if P, maybe also dispose their bearers to come to believe propositions Q such that, when combined with one's other beliefs, P counts as good evidence for Q, etc. We hope and suspect that the precise details of the functional account of belief won't matter for our purposes here. (See e.g. (Dretske, 1981, 1995; Field, 1978;

For example, we count speakers and hearers as having a lot of beliefs about—as bringing to bear a lot of information about—grammaticality, about lexical meanings, and about compositional rules. Famously, speakers of a language don't typically have explicit, straightforwardly articulable beliefs about these things. One of the reasons why linguistics is hard is that our judgments about e.g. grammaticality aren't based on an explicit articulation of a system of grammatical rules that we can access with general-purpose cognition in order to, for example, write them down on a piece of paper after a moment of introspection. But the production and comprehension of speech isn't magic—speakers of a language do, and must, have some kind of cognitive structure that underwrites their grammaticality judgments and their interpretive practice. Plausibly, that cognitive structure is implemented in the brain in a very different way than, for example, episodic memories of particular past events. We're counting both as belief, but we're happy to retreat to some more general terminology in the face of sufficiently determined objections to that usage of “belief”. The state we have in mind is, again, a very general, highly multiply-realizable state of *being in a position to bring information to bear*—one that includes paradigmatic, explicitly articulated beliefs, as well as the attitude I had toward the proposition *that no armadillo weighs more than any aircraft carrier* before Stalnaker brought it to my attention in that example in *Inquiry*, and also the attitude (sometimes called “cognizing” - see e.g. (Chomsky 1986, pp. 265-269)) that we have toward e.g. the grammatical rules of our native languages.

Note, finally, that as we're using “information” in “bringing information to bear”, there's no factivity requirement on information. It's perfectly possible for an audience to bring to bear the information (in the sense we mean) that P when P is false. Here too, we are happy to retreat to different vocabulary in the face of determined opposition to our stipulated usage.

2. “Yeah, duh”: Why goatism isn't too obviously true to be interesting

The kind of view that we're advocating isn't unprecedented in the literature. In fact, one might think it's widely presupposed in at least a significant swath of the literature. The picture of

Lewis, 1966, 1983; Schwitzgebel, 2002; Stalnaker, 1984) for some work on spelling out some of the details of this sort of functionalist account.)

linguistic communication laid out by David Lewis in “Index, Context and Content” is a prominent example:

Early in “Index, Context and Content” (Lewis, 1980), Lewis tells a story about how linguistic communication works. He says,

“Suppose (1) that you do not know whether A or B or...; and (2) that I do know; and (3) that I want you to know; and [two more conditions we don’t need to worry about just here].... Then I will be truthful and you will be trusting and thereby you will come to share my knowledge. I will find something to say whose truth depends on whether A or B or... and that I take to be true. I will say it and you will hear it. You, trusting me to be willing and able to tell the truth, will then be in a position to infer whether A or B or...”

This sounds very much like what we want to say, because it sounds like what both we and Lewis are saying is this: The way linguistic communication works is that Lewis (for example) gets us to believe that P by producing an utterance such that he’s confident we will assume that (a) it’s true only if P, and (b) it’s true. Then, from the fact of the utterance, plus our beliefs about its truth conditions together with our belief about its truth value, we can infer P. Our belief that the truth of Lewis’s utterance depends on whether P will in turn be based on a bunch of beliefs about which sounds Lewis made, about the pronunciations of various English words, and about the meanings of those words and the compositional rules governing their combination.

At this point it’s natural to worry that this sort of picture isn’t just commonly presupposed in the existing literature, but that it’s *universally* presupposed, because it’s the only way that things could possibly work. Of course that’s how updating in response to speech works. That’s how updating in response to everything works. And it’s the only way it could work. Updating just is—obviously, uncontroversially—updating with the facts of and about some observed event and inferring stuff that follows from that, together with other things you believe. (In Bayesian terms, conditionalizing on the fact of the utterance and letting the rest of your beliefs update in accordance with your priors.) And because it’s so obviously, inevitably correct, it couldn’t possibly be worthwhile pointing it out.

To this we have two responses: First, that it’s *not* the only way things could work. Second, even if it was, being explicit about it is still theoretically helpful.

The second response blurs into our response to the “so what?” reaction, and we’ll try to substantiate it in more (perhaps in painful) detail in section 4. But we can make a general version of the point right away. This Lewisian picture of linguistic communication is, as noted, congenial to our view. However, we want to emphasize that, on the informationalist view as we conceive it, the Lewisian account is an instance of a much more general communicative strategy that has nothing special to do with truth. Here is the general strategy, of which the Lewisian account is an instance: I want you to know P. So I pick an utterance to make whose F-ness depends on whether P, and such that I am confident you will trust that my utterance is F. You trust that my utterance is F. And so you come to believe that P.

Truth is a particularly interesting and frequently theoretically important substituent for F. But it’s far from the only one. Indeed it’s far from the only one that’s prominent (at least implicitly) in the literature in the philosophy of language. This is what’s happening in standard cases of conversational implicature. Betty asks Alpie to see a movie with her on Thursday night. Alpie wants to communicate that he can’t go. He says something (perhaps “I have an exam on Friday morning”) whose accordance with the cooperative principle depends on (perhaps among other things) Alpie not being able to make it to the movie. Betty trusts that Alpie’s utterance accords with the cooperative principle. And so she concludes that Alpie can’t make it to the movie. And there are many, many more features than truth and the various species of cooperativeness that could, and do, play this structural role in actual communication. (Specific examples coming, as noted, in section 4.)

In fact the underlying load-bearing phenomenon is more general still. It’s not just that there’s nothing special about truth. There’s also nothing special about utterances (understood as distinctively language-y sorts of actions). What happens in general is that I pick something to do—call it ϕ ing—whose F-ness depends on whether P, and such that I think that you will assume that my ϕ ing is F. You notice my ϕ ing, assume that my ϕ ing is F, assume that my ϕ ing’s F-ness depends on whether P, and come to believe P. Utterances are especially prominent-in-the-literature substituents for ϕ , and truth for F, but neither is terribly special with respect to the central mechanism underwriting the communication.¹¹

¹¹ Where they are liable to be terribly special is, as mentioned earlier, with respect to the kinds of background beliefs that are being exploited, and the potentially specialized mechanisms by which they are exploited, in order to support the inferences to various Ps.

Whenever the hearer has beliefs about some feature(s) the utterance has—sincerity, truth, relevance, grammaticality, vocal clarity, etc.—they can extract some additional information from the fact that the utterance occurred, and had the observable features that it had. When the speaker knows something about what the hearer believes about features the utterance has, the speaker can exploit that to deliberately convey some information to the hearer. When the hearer knows something about what the speaker believes about what the hearer believes about features of the utterance, and takes for granted that the speaker is being cooperative/helpful/etc, the hearer can extract still more information, and the Gricean communicative process can kick in.

Something that's helpful about explicitly articulating things in these terms (at least, it's helpful for us, and we hope and expect that this isn't super idiosyncratic) is that it makes it easy to see the continuity between all of the various messages extractable from, and communicable by, utterances. It also makes it easy to see that the Lewisian process of extracting information from the fact of the utterance, plus knowledge of truth-conditions and an assumption that the speaker is speaking truly, is just a particularly prominent special case of a much more general phenomenon. So while we (not surprisingly) think that the picture of linguistic communication we get from this bit of Lewis in "Index, Context and Content" is right as far as it goes, we also think it doesn't go far enough to cover the full informational and communicative richness of actual utterances. The informationalist picture is a generalization of what Lewis says that captures what Lewis is right about, but covers a lot more theoretical territory than Lewis is targeting.

So much for the general version of "it can still be helpful to explicitly articulate informationism even if everybody was already committed to it because it's the only game in town" response. We now turn to "it's not the only game in town". So let's look at an alternative model for how communication might work, and how the updates to the states of hearers and context that utterances produce could be brought about.

Here is how communication works among the Buttonheads: Each Buttonhead has an array of buttons on their forehead, wired up to their representational systems in such a way that pressing a particular button just directly causes that Buttonhead to form a particular belief. If we are Buttonheads, and you press the *giraffes outnumber chimpanzees* button on my forehead, I'll immediately and automatically form a belief that giraffes outnumber chimpanzees. (Maybe the button causes the production of a token of the language of thought sentence that expresses *that*

giraffes outnumber chimpanzees, which is then immediately deposited into my belief-box.) We could then intervene on each others' beliefs by pressing the various buttons on each others' foreheads, directly causing the target beliefs in our audiences in a way that is not mediated by any of the target's other beliefs about anything, including their beliefs about button-pressings. It's not that I regard your pressing the *giraffes outnumber chimpanzees* button as evidence that giraffes outnumber chimpanzees. Your pressing the button just directly causes me to form the belief in a way that's not evidence driven at all.

That's a cartoonish, obviously non-actual possibility. We are not, in fact, Buttonheads. But it's a close relative of some not at all cartoonish, not-obviously-non-actual possibilities. It's not hard to imagine a linguistic system that worked in an approximately Buttonhead-ish sort of way: The way updating in response to utterances works is that processing the sentence just directly causes a belief by way of the operation of a highly modularized, automatic (what (Mandelbaum, 2011, 2014) calls *ballistic*) language processing system. The informational model we are advocating, though it is compatible with modularity in general, is not a good model of this kind of system. And we certainly shouldn't rule out *a priori*—perhaps we should not rule out at all, given the total balance of available evidence—that at least some of our communicative systems work in this direct, doxastically unmediated sort of way. It's not obvious that such an account of at least some human communication isn't correct.¹² And it's especially not obvious that it couldn't have been correct.¹³ So informationalism isn't the only game in town.

3. No way: why goatism isn't too obviously false to be interesting

Here is what looks at first glance like a straightforward refutation of goatism (we think it really is a straightforward refutation of an overly-ambitious species of goatism):

¹² For some precedents for views in the immediate neighborhood, see for example (Mandelbaum, 2011, 2014) on ballistic systems of belief-formation, (Gilbert, 1991; Gilbert et al., 1990, 1993) on Spinozan/Cartesian belief formation), (Gendler, 2011) on stereotype activation, (Camp, 2013, 2017) on perspective-adoption, and (Murray, 2014) on direct update.

¹³ Because we're working with a high-altitude, massively multiply realizable, tacit-belief-allowing Stalnakerian picture of belief, it's also not obvious that this kind of forehead-button-ish picture *is* correct. Given our permissive notion of belief, we can't move from the (clearly correct) observation that lots of updating in response to people's utterances isn't *consciously* doxastically mediated to the conclusion that it's not doxastically mediated at all.

Goatism can only deliver doxastic/presuppositional updates. But lots of the features of context that get updated are not, or not entirely, doxastic (QUDs, to-do lists, planning states, etc.). Updates to the non-doxastic features of audiences/context are not going to be goat updates.

We agree with all of this except the part that says that goatism can only deliver doxastic or presuppositional updates. What's true is that goatism can only *directly* deliver doxastic or presuppositional updates. But this doesn't actually make trouble for the goatist picture, once it's stated somewhat more carefully than we've been doing so far.

We think that, for example, the update that happens when a question is successfully asked - adding a question to the stack of questions under discussion - is reducible to the attitudes of the conversational participants, but isn't reducible to their beliefs. Having a question on the QUD stack is no doubt in part a matter of the group's collective beliefs, but it's also (very plausibly) a matter of collective non-doxastic attitudes as well - intentions and/or plans, for example, or maybe irreducible inquisitive attitudes.¹⁴ Still, the way questions get added to the QUD stack is (we propose) *by way of* a doxastic update, and the right story about the doxastic update is goatist.

Tim asks, "where's the beer?". Andy, bringing to bear some background information - including some facts about English phonology, syntax, and semantics - forms the belief that Tim intends to add the question of where the beer is (let's call this question Q) to their conversational agenda. Andy also believes that Tim is a competent and well-intentioned conversation partner, and so is confident that if Tim wants to add Q to the agenda, it's a good idea to add Q to the agenda. So Andy comes to believe that it would be a good idea to do his bit toward adding Q to the agenda. (And also forms an expectation that Tim will do his bit, that Tim will expect Andy to do his bit, etc.) And then Andy, as a result of forming the belief that it would be good to add Q to the agenda, does his not-entirely-doxastic bit toward adding Q to the agenda. And so (assuming Tim does his bit as well) Q goes on the stack. A psychologically reducible, but not doxastically reducible, update to the conversational context happens, and the way it happens is as a downstream consequence of a goat-friendly doxastic update.

¹⁴ See e.g. ((Friedman, 2013; Roberts, 2004, 2012)

This generalizes. Stipulate a linguistic practice where the conventional response to use of a sentence S is for audiences to ϕ , for some non-doxastic ϕ .¹⁵ Then we maintain that the way we actually wind up ϕ ing in response to uses of S is by each of us recognizing that the speaker produced S , recognizing that the speaker's production of S calls for us to ϕ , assuming that if the speaker is calling for us to ϕ , then it's a good idea to ϕ , concluding that it's a good idea to ϕ , and on that basis going ahead and ϕ ing. Whatever non-doxastic updates you want to build into your story are then brought on board as downstream consequences of goat update.

Alternatively, we can characterize the conventional responses to speech acts in the first instance in terms of updates to the conversational context, rather than to the states of particular audience members. In that case we let ϕ be something of the form, *doing one's individual bit toward producing update U to the conversational context*. This will allow us to accommodate any context update that's reducible to facts about the mental states of participants in the conversation - presupposition, collective intention, coordination on a class of eligible planning states or probability distributions, etc.

So our slogan, more carefully stated, is: All update is goat update, or a downstream consequence of goat update.

That's still a substantive claim. The causal mechanisms that give rise to non-doxastic updates could be doxastically unmediated. It could be that asking a question causes collective intention formation (etc.) directly, in a Buttonhead-ish way, not by way of producing beliefs that it would be a good idea to form the relevant intention(s) (etc.). But it's plausible, we think, that the psychological updates that either are, or constitute, or bring about, the updates to the states of audiences and/or contexts as a result of linguistic communication are in fact uniformly doxastically mediated. In particular, they're mediated by beliefs about which sentences were produced and/or which speech acts were performed, about the kinds of responses that are conventionally called for in response to the production of those sentences and/or speech acts (and/or which the response(s) that the speaker intended to bring about), and about the reliability of the speaker's performance of an action that conventionally calls for (and/or is intended to produce) response R as an indication of the advisability of R .

¹⁵ The further generalization to sentences that call for different responses in different contexts of use is straightforward but complicates presentation. We trust that the reader who knows enough to be concerned also knows enough to be able to add "in context c " at the appropriate points.

The general strategy here is to allow the informationalist to fold all of the communicative effects that aren't themselves directly informational into their account by taking them to be downstream consequences of informational updates, mediated by (typically tacit) knowledge of the facts about linguistic meaning that connect the sentences with the relevant non-doxastic updates. Here is a shorthand/metaphor for this move, suggested by Chris Barker (pc): Think of the association of sentences (in context) with update effects as a sort of rulebook. Let the rulebook associate sentences (in context) with whatever kinds of updates you like, doxastic or otherwise. The goatist can fold all of these updates into their account by allowing the goat to eat the rulebook. Goatists can explain all kinds of non-doxastic responses to utterances by attributing them to (in part) beliefs about what the rulebook says—beliefs about the rules (norms, conventions, etc.) that determine which non-doxastic responses are called for in response to which (kinds of) utterances.¹⁶

4. “So what?”: Why framing things in explicitly goatist terms is helpful

We think framing explanations of linguistic communication in these terms is helpful with respect to a surprising range of issues. We'll pick a few of what we think are the headliners to discuss here.

The most obvious one is theoretical unification. The goatist, informationalist approach pulls all the cases from our introductory example under a common theoretical umbrella, and makes it easy to see how they're related to each other. And approaching all of these different types of communication in an informationalist framework is, we have found, helpful in putting a spotlight on some phenomena that are otherwise easy to miss, and renders unsurprising some phenomena that might initially have seemed odd or *recherché*.

¹⁶ One kind of view of the ways in which speech acts change conversational contexts that this rulebook-eating maneuver can't assimilate is the kind of view according to which the context-changes produced by speech acts are not just non-doxastic, but entirely non-psychological. (Or on which the psychological changes are downstream consequences of more-basic changes that are not at all psychological.) For example, views according to which the essential effects of speech acts are changes to the *normative* facts. So views according to which, for example, what promises are fundamentally in the business of doing is producing obligations, not beliefs (or expectations, or whatever), and the essential effect of an assertion is to saddle the speaker with a distinctive sort of *commitment*, and any effects on belief or presupposition are downstream effects of this more-basic normative effect. (See e.g. (Camp, 2018; Caponetto, 2020; Caponetto & Cepollaro, 2022; Kukla, 2014; Kukla & Lance, 2009; Langton, 2017), and probably also (Brandom, 1983, 1994; Rescorla, 2009). Two informationalist responses. Aggressive goatism: resist these views, either denying the normative effects really occur, or reversing the Euthyphro order so the normative effects are consequences of doxastic updates. Concessive goatism: Goatism as an account of psychological and downstream-of-psychology updates. (Which is still a lot!)

We'll illustrate by running through some examples, calling back to the example of Alphonse's utterance at the Central APA that we opened with. The whole enormous, diverse array of potential things Betty could learn, and Alphonse could communicate (of which 1-26 are just a small sample), is explained in a common framework whose only moving parts are variation in the information brought to bear by hearers and variation in the speaker's beliefs (hopes, plans, etc.) about the information hearers will bring to bear, on the observable facts of and about the utterance. Along the way we will note some places where informationalism sheds helpful light on some interesting communicative phenomena.

4.1 Paradigmatic literal update

In the paradigmatic case of a hearer updating with the literal content of a speaker's utterance, the speaker produces a sentence *S*, which expresses (in context) a proposition *P*, and the hearer updates with *P*. In our opening example, this is Betty's update with

(1) Carrie will be in the ballroom of the Palmer House Hilton at 1:05 pm.

On the informationalist picture we're advocating, that's what happens when the hearer updates with the fact of the utterance, while holding fixed a specific array of background beliefs about it. So for example, if Betty knows which sounds Alphonse made, knows which English sentence those sounds are a pronunciation of, knows what all the words in the sentence mean, knows enough about the context-dependence-resolving facts to resolve any context-dependence, and assumes that Alphonse is speaking truly in English, then Betty will update her beliefs with the proposition Alphonse's utterance expressed. That's the paradigmatic case of updating with literal content of a speaker's utterance.

At least, that's the paradigmatic case of hearers updating their own doxastic states with the literal content of the speaker's utterance. We can generalize to collective updates. Where *P* is the literal content of Alphonse's utterance: If Betty thinks that all the other parties to the conversation are making similar assumptions about Alphonse's utterance, and so they're updating with *P* as well, and also thinks that all parties to the conversation are assuming that all parties to the conversation are so updating (etc.), then Betty will wind up *presupposing* *P*. And if all the

other parties to the conversation really are making similar assumptions, then we'll wind up with P entering the conversational common ground.

It's also easy to characterize cases where we get collective updates that fall short of full-dress, iterated-all-the-way-up presupposition.

If the participants in the conversation are making somewhat weaker assumptions - ones that are strong enough to generate some but not all of the layers of iterated belief required for full presupposition - then we'll get P assuming a status somewhat weaker than presupposition. (For example, P might wind up universally believed, and universally believed to be universally believed, but not universally believed to be universally believed to be universally believed.)

If some but not all of the participants in a conversation are making the required assumptions about meaning, truth, and each others' assumptions, then we'll get P becoming common ground among a proper subset of the conversational participants. This is what will happen, for example, in the case where Dan's zoom video is frozen but we think maybe the audio is getting through to him, though we're not sure. Then the content of Tim's assertion will become common ground between Andy and Tim, we'll all be sure it's *not* common ground among all three of us (because even if Dan did get the message, Tim and Andy suspend judgment on whether he got it, and Dan knows this), but we'll leave open the possibility that it's universally believed among the three of us.¹⁷

That's how the informational picture accounts for paradigmatic cases of acceptance of literal assertion. Well, almost. We complicate the story about literal acceptance a bit with a qualification, an elaboration, and a note, before transitioning to our next category of update.

Qualification: this is how the informational imperialist accounts for paradigmatic cases of literal assertion when we're just thinking about belief, rather than acceptance.¹⁸ The story about acceptance will be a rulebook-eating story (see section 3)—acceptance happens when, and because, hearers bring to bear enough information to recognize that they're being called upon to accept some proposition P, trust that the speaker has good judgment about what it's a good idea to accept in their conversation, and so accepts P on the basis of their newly-acquired belief that it's a good idea to accept P (and perhaps also that the speaker, and other parties to the

¹⁷ This case is a variation on the kinds of cases that feature prominently in (D. W. Harris, 2020).

¹⁸ See for example (Camp, 2018; Stalnaker, 1978, 2002; Yalcin, 2007) on acceptance vs. belief.

conversation, are expecting them to accept P). The story about addition to the acceptance-based common ground will be as above.

Elaboration: Above, update is just the result of assuming the speaker is speaking truly in English, plus the facts about English, utterance, and context that fix which proposition was expressed in English in the utterance's context. We might (stronger—we certainly will) also want a role for recognition of the speaker's communicative intention. We can get that from a richer array of background assumptions. If Betty knows the facts she needs in order to know that Alpie's utterance expressed P, and also thinks that Alpie produced the utterance deliberately, with the intention of getting Betty to believe the proposition expressed, then Betty will extract that Alpie intended, by producing S, to communicate P. If she also assumes that whatever Alpie is intending to communicate is true (or, weaker, that if Alpie is intending to communicate P, then P), then she again gets to P, this time by way of a recognition of Alpie's intention to communicate P.¹⁹

There are a lot of background beliefs required for successful communication here. One might be worried it's too much extra belief. But we remind the reader that we are, throughout, assuming an account of belief that allows for a lot of tacit belief that guides action without explicit or conscious articulation by the believer. And given that undemanding, high-altitude, mechanism-neutral notion of *belief*, all the requirements seem plausible—if Betty *didn't* believe (at least tacitly) that Alpie was speaking English, etc., it would be unreasonable for her to update with the proposition actually expressed. Also, an account like ours according to which paradigmatic cases of literal update depend on all of this rich background of belief will predict that we should see a variety of kinds of communicative failure, and/or alternative communicated contents, in cases where some of the relevant beliefs go missing. The fact that we do in fact see such phenomena (we'll draw attention to a number of them shortly), and that an informational account is well-positioned to theorize about them in an enlightening way, is a point in favor of it.

¹⁹ Footnote elaborating on elaboration: The assumption that the proposition expressed by the sentence is true, and the assumption that the proposition that's the object of the speaker's communicative intention is true, can come apart, and we can get different updates depending on what's held fixed. Karen could think that the sentence Devin produced is true (in English, in context) even if she doesn't think that Dan has a communicative intention about the proposition the sentence expresses (in English, in context). (Imagine for example a case in which Superman, using his super hearing, overhears Lex Luthor soliloquizing about his plans when he thinks no one can hear him.) And Liz can think that the object of Chris's communicative intention is true even if she doesn't think that the sentence Chris produced semantically expressed a truth. (One class of cases where we may want to say this is cases of hearers compensating for speakers' malapropisms (Davidson, 1986) Another is (at least some) cases of non-literal speech. Another is the class of *partial information updates* which we will go on about at some length below.)

Transition: There are lots of cases in which we see updates that are results of bringing to bear packages of background beliefs other than the paradigmatic one that delivers the truth of the proposition that's literally expressed. One prominent type of case is the one where the package of information brought to bear on the utterance is weaker than one that would support update with the proposition expressed. Many of these are cases of communicative failure and breakdown due to ignorance or error by the speaker, the hearer, or both. But as we'll see, not all of them are cases of failure or breakdown. One thing we think is helpful about thinking about updates in response to assertion in a goattist, informationalist way is that it makes it easier to see what the standard sort of successful communication depends on, and to diagnose and theorize about what happens in various kinds of breakdown, when some of the assumptions that support successful communication in paradigmatic cases go missing, as well as cases of communicative success, despite the fact that some of those assumptions go missing. It's also helpful for thinking about the role of assumptions *other* than truth in supporting the sorts of communicative effects our utterances to one another actually have.

4.2 Updates that don't rely on assuming truth

At this point it will be useful to introduce a notational convention that we'll make use of in what follows. We'll use expressions of the form, $[A + B + C\dots]$ to talk about the specific pieces of information, or types of information, that hearers bring to bear. So for example we'll say that standard cases of updating with the literal content of assertion are supported by the hearer bringing to bear [utterance + proposition expressed + truth], meaning that updating with the literal content of an assertion is supported by hearer's bringing to bear (a) the information that the utterance occurred and had the observable properties it had, (b) the information that the utterance expressed the particular proposition that it did, and (c) the assumption that the proposition expressed by the utterance is true. Depending on what's most presentationally helpful, we'll sometimes talk about hearers bringing to bear $[utterance + A + B]$, and sometimes about hearers bringing $[A + B]$ to bear on the utterance. In using this convention, we don't have anything fancy or technical in mind. It's just that we are going to have occasion to talk a *lot* in what follows about different packages of information that a hearer might bring to bear in order to extract information from an utterance, and this sometimes makes presentation easier.

One straightforward sort of case of updating based on information weaker than what would support update with the literal content of the assertion is the case where the hearer has enough information to determine which proposition the utterance expressed, but isn't confident that the utterance is true. The most straightforward cases of accepting that the proposition the speaker expressed with their utterance is true are based (in part) on the hearer being confident that the speaker is sincere, and that their beliefs about the subject matter are correct. (In the paradigmatic case, confidence in *truth* is based on confidence in *sincerity* and *competence*.²⁰ In other words, a hearer's bringing to bear [proposition expressed + truth] in order to extract an update from an utterance is, in many cases, the product of bringing to bear [proposition expressed + sincerity + competence].)²¹ (One more terminological stipulation: we'll use "proposition expressed" to mean "proposition semantically expressed"—that is, the proposition that the correct semantics for the speaker's language associates with the uttered sentence in the speaker's context.) The most straightforward cases of updating on the basis of something weaker than truth are cases where the hearer isn't confident of the speaker's sincerity, or of the accuracy of the speaker's beliefs, or both.

On an informational picture it's easy to see that, and how, the hearer's information about the sentence produced and its meaning will still support a number of updates, if not with the truth of proposition in fact expressed, then with (for example) some propositions about the speaker's attitude(s) toward the literal content of their utterance. If from [proposition expressed + sincerity + correctness] we could get *P*, then from [proposition expressed + sincerity] we can get *that the speaker believes P*. From [proposition expressed + motivational premises other than sincerity], we'll often be able to get things like *that the speaker wants me to believe that P*, *that the speaker has interests that they think will be well served by my believing that P*, and/or *that the speaker wants me to think that they believe P*.

In general, there will often be a fair bit of information that's extractable from [proposition expressed + something that falls short of securing truth], where the something else is frequently one of the components of a cluster that would, in more favorable cases, get us to truth. (For

²⁰ See for example (E. Fricker, 1994; M. Fricker, 2007; Hardwig, 1985; Lackey, 1999, 2006, 2008).

²¹ The information about which proposition is expressed can also usefully be disaggregated in various ways - for example, to [meanings of sub-sentential expressions + context-sensitivity-resolving facts + principles of composition].

example sincerity, or justification, or having compelling evidence.)²² This is what's happening with (2) and (3) from our original example:

2) Alphie believes that Carrie will be in the ballroom of the Palmer House Hilton at 1:05pm.

3) Alphie has compelling evidence that Carrie will be in the ballroom of the Palmer House Hilton at 1:05pm.

Notice that these updates don't compete with the update with the proposition in fact expressed. In very many cases, Betty will (predictably) update with all of (1), (2), and (3), and all three will become part of the common ground in the conversation in which Alphie's utterance takes place. The informational resources required to get to each are different, but Betty will (in many cases) have all three required packages of information available. (In many versions of the case, for example, the resources Betty brings to bear to extract (2) will be a proper subset of the ones she uses to extract (1).) And it is perfectly possible—indeed it's quite common—for a speaker to predict and intend that their audience update with (the analogues of) all of (1), (2), and (3).

This is an instance of a general point that we'll emphasize a lot in what follows—approaching communication in a goatist framework, there's no reason to expect that utterances will be associated with just a single communicative effect, or that there will be exactly one proposition that's a target of a speaker's communicative intention. Much less that the proposition semantically expressed in context captures the whole of the utterance's communicative effects, or is the sole eligible target for a speaker's communicative intention. Indeed, there's plenty of reason to expect that speakers will ubiquitously have more than one communicative intention with their utterances. We count this as a virtue of an informationalist approach, because it's clearly true that speakers very often have multiple communicative aims with their utterances. Informationalism makes it easy to see why this is to be expected, and provides a framework that makes it easy to theorize about them in all of their rich multiplicity.

²² There's an interesting (to us, at least) point of connection to some of the kinds of cases that Jennifer Lackey uses to argue that testimonial knowledge doesn't depend on speaker knowledge. (Lackey, 2008). Some Lackey cases will deliver update with literal content but not update with the proposition that the speaker believes literal content, from e.g. [meaning fixing + j-tracking + J->T], because the package of beliefs that secures the truth of the utterance isn't the standard sincerity + correctness package.

Relatedly, the informationalist framework also gives us the resources to clearly distinguish a variety of updates, based on differences in the sorts of communicative intentions or expectations, if any, that the speaker has about them.

There are the *extractable* updates—the updates such that the hearer does, in fact, have the information required to extract them, whether the speaker knows this or not. This is an extremely unselective category. There will, more or less always, be very, very many of these.

There are the *predictable* updates—the updates such that the speaker expects that the hearer will bring to bear the information required to extract them. This is a somewhat more selective category. There will be, typically, some update-supporting information that hearers have that speakers don't know about, and so there will be some updates that the hearer will be able to extract though the speaker won't be in a position to predict this. (There will also be some mis-prediction, in which speakers anticipate that audiences will have some information that they don't actually have.) But there will still be quite a lot of these associated with pretty much every utterance.

There are the *predictable, intended* updates—the updates such that the speaker both expects *and intends* that the hearer will bring to bear the information required to extract them. Here we get a constraint that will thin the herd quite a bit - very many updates will be predictable but not intended. (Obvious examples are, in typical cases, a bunch of predictably extractable updates about laryngitis, vocal range, accents and geographical and/or cultural origin, linguistic competence, etc.. Also often the updates about the speaker's belief state or evidential situation. But there will be a lot of variation across cases, and anything that can be a predictable update seems like it could, in the right kind of circumstance, also be intended.)

Then there are the updates that are the objects of Gricean communicative intentions. This is quite a selective group. The speaker needs to have some pretty fancy intentions about P in order for it to fall into this category, and even many of the updates that the speaker intends their audience to pick up on will be excluded, because the speaker doesn't intend the update to proceed by way of the audience's recognition of their intention. But there's still no reason to expect uniqueness here. Most obviously, in paradigmatic cases of conversational implicature the speaker has a Gricean intention about what's literally said and also about what's implicated. Implicature is not special in this respect. It's perfectly possible to have Gricean intentions toward the literal content of one's utterance, toward one or more implicatures, toward an update based

on an honorific, toward an update about one's class background, and also toward an update about the impressive extent of one's vocabulary, for example. And we should not be surprised to see a mix of speaker predictions and intentions of varying degrees of complexity associated with a single utterance. It shouldn't be surprising, for example, to find a speaker (merely) predicting that their audience will update with P, predicting and intending that they update with Q, and having a full Gricean communicative intention toward R.

Another virtue of the informationalist approach is that it suggests (we think correctly) that there's no terribly sharp joint in nature that singles out the updates that are the objects of full Gricean communicative intentions as the one and only thing that it could ever make sense to talk about when we're talking about communication or meaning, or even as ones that have a uniquely privileged role in talk about communication or meaning.²³ All of the various points along this spectrum are likely to be interesting and useful to theorize about. So we shouldn't be content with the sort of Gricean taxonomy of meaning that just has two buckets—one for natural meaning, and one for non-natural meaning. The various points along the way, where speakers have variously complex intentions that fall short of the full, three-layers-deep-in-metacognition kinds of intentions required to support Gricean meaning_{NN}, are plenty interesting too.

This approach to the updates in (2) and (3) delivers a straightforward account of individual update on the basis of a package of background beliefs that doesn't include utterance truth. There's also the usual straightforward extension to collective update. When everyone expects everyone else to bring to bear a package of information that includes content-determining information, but not truth (and expects everyone else to expect everyone else to bring that package to bear, etc.), we'll get the relevant propositions about (for example) the speaker's attitudes toward, or evidential situation with respect to, the proposition expressed being added to the conversational common ground.

In the case where the speaker was anticipating that the audience would assume truth and update with the content expressed, this is liable to result in a defective context, where the speaker presupposes the proposition expressed by the utterance and their audience doesn't. But there will also be plenty of cases where the speaker won't be assuming (especially after they see their audience's reaction to their utterance) that their audience has accepted the truth of the proposition they expressed, and so we'll wind up with a non-defective context that's been updated with the

²³ Compare (Armstrong, 2018, 2023)

information about (for example) the speaker's attitude, but not with the truth of the proposition expressed.

4.2 Updates based on assuming more than truth - standard implicatures

We observed above that straightforward cases of successful literal assertion can be thought of as updating on the fact of the utterance, against the backdrop of assumptions about the linguistic meanings of the sounds uttered, assumptions about the context-dependence-resolving facts, plus the assumption that the speaker is speaking truly. We also noted that it's also possible to extract (and to predictably and deliberately communicate) a lot of information from an utterance on the basis of weaker background assumptions - in particular, without relying on the assumption that it's true.

It's also worth looking at what happens when hearers bring to bear (and potentially, when speakers and hearers are relying on each other to bring to bear) stronger background assumptions than the ones that support standard cases of literal update. One thing that happens is that, when we add the further assumption that the speaker's utterance is *cooperative* in Grice's (Grice, 1975) sense, we get the standard sorts of Gricean implicatures. (Terminological note: In the official Gricean taxonomy, truthfulness is part of cooperativeness. We will deviate from official Gricean usage by using "cooperativeness", to mean what scrupulous Griceans would mean by "cooperativeness with respect to considerations other than Quality," since for our purposes it's helpful to separate *cooperativeness*, as we'll be talking about it, from *truth*.)

Take Alphie's reply to Betty's invitation to see a movie on Thursday night: "I have an exam on Friday morning." Betty updates not only with the proposition that Alphie has an exam on Friday morning, but also with the proposition that he is thereby unavailable to see a movie on Thursday night. The implicature here is smoothly assimilated to the informationalist picture. From the fact that Alphie made certain noises, plus some facts about English pronunciation and grammar, we get *that Alphie produced the English sentence "I have an exam on Friday morning"*. From that plus assumptions about the meaning of the expressions and some relevant aspects of the context, we get *that Alphie's utterance expressed the proposition that he has an exam on Friday morning*. From all that plus truth, we get *that Alphie has an exam on Friday morning*. And from *that*, plus cooperativeness, we get that Alphie can't make it to the movie.

Put another way, Alpie produces an utterance that Betty assumes to have certain features—not just various semantic features, and not just truth, but also cooperativeness. Given the specifics of their context, Betty knows that the cooperativeness (and specifically, relevance) of Alpie’s utterance depends on his not being available to see the movie on Thursday, just as its truth depends on his having an exam on Friday morning. So from [proposition expressed + truth + cooperativeness], along with her assumptions about what cooperativeness demands in their context, Betty can extract not just that Alpie has an exam, but also that he can’t make the movie. Alpie is confident that Betty will trust that his utterance is not just truthful but also cooperative (and specifically, relevant). And so he can use his utterance to intentionally communicate not just that he has an exam, but also that he can’t see the movie.

Nothing in this story goes beyond the usual Gricean account of a conversational implicature like Alpie’s. What we want to observe is that this mechanism for communicating information is distinctive *only* in the background assumptions involved, and not in any deeper way, from the communicative mechanisms underwriting other sorts of updates. We can learn a lot from the fact that somebody said something, in addition to the literal content of what they said. Conversational implicature is what happens when the assumptions relevant to the hearer’s update include not just truth, but cooperativeness, or cooperativeness in a particular respect.

This is what happens with items (4)-(7) from our opening example:

- 4) Whether Carrie will be in the ballroom of the Palmer House Hilton at 1:05pm is relevant for the purposes of Alpie and Betty’s conversation.
- 5) Alpie believes that Betty didn’t already know that Carrie will be in the ballroom of the Palmer House Hilton at 1:05pm.
- 6) Carrie is not presently in the ballroom of the Palmer House Hilton.
- 7) Carrie will be able to give her talk.

In fact more than just [proposition expressed + truth + cooperativeness] is probably required to get to (7). Betty probably also needs additional background information about the prevailing norms about how late it’s ok to start talks, for example. Very many standard cases of conversational implicature are like this—cases like (4)-(6), which are plausibly *just* based on the assumption of cooperativeness, are probably the exception. Consider for example 8-10:

- 8) Carrie will be late for her talk.

9) Carrie is not in San Francisco.

10) Betty should stall, rather than canceling the session and asking everybody to leave the ballroom.

All of these require, in addition to [proposition expressed + truth + cooperativeness], some extra information—for (8), about the scheduled starting time of the talk, for (9), about geography and about what kinds of transportation options are likely to be available, and for (10), about the sorts of goals and contingency plans Alphonse, Betty, Carrie, and probably also the audience, are working with and open to.

Just as previously we got different, and still interesting, updates from assuming not truth but sincerity, we also get different, and still interesting, updates when we assume not cooperativeness, but attempted cooperativeness. Getting to (4) and (5), for example, doesn't require assuming that what Alphonse said is *true*—just that it's relevant, or a useful increment of additional information. And now consider (11), extraction of which doesn't depend on *relevance* itself, but on *attempted relevance*:

11) Alphonse believes that whether Carrie will be in the ballroom of the Palmer House Hilton at 1:05pm is relevant for the purposes of Alphonse and Betty's conversation.

This provides another example of the point we promised we'd return to in the discussion of (2) and (3). There's clearly no competition between (4) and (11) as potential things to update either a hearer's doxastic state or the conversational common ground with. It's perfectly possible, and indeed likely to be very common, to update with both. It will be rare and weird for a hearer to assume relevance without assuming attempted relevance, and so it will be rare and weird to update with (4) but not (11). (Cases of the reverse are easy to imagine, though—for example, cases where the speaker has lost the thread of the conversation and failed to track a change in topic). It's also worth noting that, since there's no entailment in either direction, there's no redundancy here. It's not the case that once you've got one of the updates, the other is informationally idle. (This is another point that we will be returning to in what follows.)

In general, whenever we assume the speaker was being cooperative we can extract extra information from the fact of the utterance, of exactly the type that Grice drew attention to. And if the speaker assumes that their audience will assume they're being cooperative, speakers can exploit that to deliberately communicate the extra information. But what's special about the

cooperative principle isn't that cooperativity is the only auxiliary assumption that can play this role. It's that, like truth, cooperativity is an example of an assumption that we're very often in a position to make, and that speakers are therefore often in a position to exploit in service of their communicative ends. In cases where we're in the market for extra bits of communicative force to attach to a speech act, cooperativity can often play this role. But it's worth noting that, and it's a virtue of the informationalist picture that it makes it very easy to see that, there are ever so many other assumptions we could make about somebody's utterance, and all of them will allow us to extract additional information. And all of them could be such that we assume them, the speaker assumes we assume them (etc.), and so all of them could be the basis of either private information extraction by hearers, or of deliberate exploitation for purposes of message-layering by speakers.

4.3 Updating on the basis of assumptions that have nothing much to do with truth

We're going to skip past numbers 12-19 for the moment, since they all go together in a section that's quite long, so we'll do the quicker stuff first.

One kind of update that's based on holding fixed a feature of the utterance that's got nothing much to do with truth—and indeed, nothing much to do with truth-conditions—is the sort of update associated with expressions like honorifics, which make important contributions to communicative force without contributing to truth-conditions, or whose contributions to communicative force outrun (in that they can't be fully explained in terms of) their contributions to truth-conditions.²⁴ Honorifics, though, are just one example of expressions that contribute to linguistic communication without necessarily impacting truth conditions. Many instances of code-switching and style-shifting will work like this, along with discourse markers, tag questions, and conversational-flow-regulating expressions like “so,” “well,” “you know,” “actually,” “I mean,” “alright?”. Depending on your view of evidentiality, or the expressions in question, evidentials may also be an example of this.

These are all cases where it looks like what happens is that some aspect of linguistic meaning determines a kind of felicitous-use condition that isn't a truth condition.²⁵

²⁴ See for example (Brown & Levinson, 1978; Camp & Nowak, n.d.; Kaplan, ms; Predelli, 2013).

²⁵ See for example (Camp & Nowak, n.d.; Kaplan, ms; Predelli, 2013)

If Andy walks up to President Biden and says, “du bist der Präsident”, that’s inappropriate (for a number of reasons, but among them is Andy’s use of the familiar “du”), even though it’s not false. That is, there’s a condition on felicitous use that’s not satisfied, even though the truth conditions are satisfied. So these kinds of updates (about familiarity, or believed familiarity, or intentions to establish familiarity, for example) are going to be, on an informationalist view, the result of updating with the fact of the utterance, while holding fixed that the utterance has some kind of appropriateness-ish feature that’s not tied up with the utterance’s truth-conditions. This is what’s happening with 20-22:

- 20) Alphie is in a position of lower social standing or authority than Betty.
- 21) Alphie takes himself to be in a position of lower social standing or authority than Betty.
- 22) Alphie is concerned to explicitly signal that he takes himself to be in a position of lower social standing or authority than Betty.

We can get (20) from information about the appropriate-use-conditions for “Professor”, plus the assumption that those conditions are satisfied (Let’s write, [appropriateness-conditions + appropriateness].) We can get (21) from [appropriateness-conditions + attempted appropriateness]. We can get (22) from the appropriateness-conditions, plus a pretty sophisticated (but very often available) package of additional assumptions about the information Alphie expects his audience to bring to bear.²⁶

Again, there’s no competition between these updates and the ones based on truth-conditions plus an assumption of truth, either as pieces of information for audiences to extract, or as objects of speakers’ communicative intentions (of varying degrees of complexity).

Another important class of cases is ones where the information brought to bear in order to generate the update doesn’t have much of anything to do with any kind of linguistic meaning, truth-conditional or otherwise. Updates 23-26 from our initial example are like this:

- 23) Alphie speaks English.
- 24) Alphie believes that Betty speaks English.

²⁶ Something like: [appropriateness conditions + Alphie wouldn’t use an expression with lower-status-than-addressee appropriateness conditions unless he was concerned to explicitly signal that he takes himself to be lower status than his addressee], where the second component is certainly going to be decomposable into interesting subcomponents.

- 25) Alpie is not from Australia.
- 26) Alpie does not have laryngitis.

The first two of these—(23) and (24)—are perhaps borderline cases as to whether they're based on "anything much to do with any kind of linguistic meaning". They depend on recognition of the language being deployed in the utterance, and maybe that has got something to do with some kind of linguistic meaning. But it doesn't have anything much to do with any of the specifics of the linguistic meanings of the particular expressions used.

We can get (23) from, for example [Alpie's utterance is an utterance of an English sentence + Alpie wouldn't be uttering English sentences if he didn't speak English]. (Note that while this second assumption is one we'll often make, it's not one we'll always make. If we think that maybe Alpie is reciting lines from a poem or lyrics from a song in a language he doesn't speak, or trying to fool the military police by producing the one sentence of English that he knows how to pronounce, or asking for a translation, we won't make this assumption.)

And we can get (24) from [Alpie's utterance is an utterance of an English sentence + Alpie wouldn't be uttering English sentences if he didn't think Betty spoke English]. This last assumption is again one that we'll very often, but not always, make when updating in response to people's utterances. (We leave the generation of cases as an exercise for the reader.)

The last two—(25) and (26)—are clearer cases where the update is completely independent of any assumptions about meaning. These we get from some facts about what the utterance sounds like, plus an assumption that the way the utterance sounds is reflective, in the usual way, of some fact about the condition of the speaker's vocal apparatus, or where they're from. (They're not faking their laryngitic status, or putting on an accent.) This is also what's happening in cases where we learn whether someone is from Canada based on how they pronounce "pasta". (Here too it's a bit tricky to separate *completely* from meaning—both (25) and the "pasta" update are plausibly at least loosely tied up with something about meaning, since in order to attribute the accent to the speaker, we've got to take them to be pronouncing a particular word or words. And so (26)'s full non-linguisticness is probably usefully separated from even (25).)

Another relevantly similar kind of case is when Tim successfully communicates to his family that he's home from his trip to the grocery store by calling out "hello!" as he enters the

house, or when the skier trapped under the avalanche signals their location to the rescue team by shouting “help!”. In all of these cases, the relevant update is extracted on the basis of facts of and about the utterance (who produced it, where it came from, what it sounded like), together with some not-even-truth-adjacent features. And while often these sorts of updates will be unintended side-effects of producing an utterance with a very different communicative aim, that needn’t, and won’t always, be the case. There’s no barrier to these sorts of not-at-all-about-truth updates being among the things that the speaker intends to communicate to their audience with the utterance, or even to their being the primary communicative purpose that motivates the utterance. (Tim’s ‘hello’, and many instances of the skier’s ‘help’, are like this, for example. But the same could go for any of (23)-(26) - it just takes a bit more creativity to generate the case.)²⁷

We would like to re-emphasize at this point something that these examples illustrate—that there’s no uniqueness requirement on being an object of a speaker’s communicative intention. The skier could perfectly well intend *both* to communicate that they need help, and *also* communicate their location, by means of their shout of ‘help!’. (Indeed, that seems like it’s probably the usual case.) And on any natural elaboration of the case, Alphie is sure to intend to communicate more than one of the items on the list of (1)-(26), and likely a number of other things that don’t appear on our list. One thing that we think is helpful about a goatist approach to update and communication is that it makes it very easy to see both that there’s no need to build the whole communicative force of an utterance into the proposition expressed, and that there’s no need to single out a unique proposition associated with the utterance as *the* proposition communicated, or even as *the* object of the speaker’s communicative intentions. (John Perry and co-authors²⁸ warn against these mistakes, under somewhat different modes of presentation but for very similar reasons (actually maybe for the same reasons under different modes of presentation). We’ll have more to say about Perry’s discussions of these sorts of issues in the next subsection.)

4.4 Updating on the basis of partial information about meaning

²⁷ For some relevantly similar cases and discussion, see (Ramírez, 2010).

²⁸ (Barwise & Perry, 1983; Israel & Perry, 1990; Korta & Perry, 2011; Perry, 1988, 2001)

Many of the updates discussed so far have been ones that depend on hearers having full semantic information—having all the information that they need in order to determine the proposition expressed by the utterance. (As well as any semantically-encoded but non-alethic appropriate-use conditions.) Being in a position to do that relies on bringing quite a lot of background information to bear. The hearer needs to have correctly perceived which sounds were made, and also to be making the right assumptions about pronunciations of words, about linguistic meaning, and about the resolution of context dependence. That’s a demanding condition that actual hearers don’t always satisfy. In many actual cases, a hearer won’t be in a position to bring to bear all of the information that they’d need in order to determine the exact proposition that’s in fact expressed by the speaker’s utterance.

Sometimes a hearer will be ignorant of some feature of context on which the resolution of context-dependence depends. Sometimes a hearer will be ignorant of the standing meanings of some word(s) in the utterance. Sometimes they won’t have correctly perceived which sounds the speaker made. Sometimes they’ll be uncertain about how to resolve a lexical or structural ambiguity. Less frequently, the hearer won’t be sure about which language the speaker is speaking. (See e.g. Davidson’s (Davidson, 1968) example of “Empedocles leaped”/“Empedocles leibt”, and Korta and Perry’s (Korta & Perry, 2011) “Nina is John”/“Ni naiz John”.) In all of those cases, even if the hearer takes for granted that the speaker is speaking truly, they won’t be in a position to update with the proposition that the speaker’s utterance in fact expressed. That’s because, for one reason or another, they aren’t in a position to complete the compositional process to arrive at the proposition expressed, because there are some points where they don’t know which objects to plug into the compositional machinery. But they will nonetheless be in a position to extract quite a lot of potentially useful information from the fact of the utterance, together with some other background assumptions - frequently and prominently, an assumption of truth.

A simple case: Betty trusts that Alpie is speaking truly, but doesn’t know enough about the conversational situation to resolve the context-dependence of “she”. So she can’t get from the fact of Alpie’s utterance to (1). She can, however, get to (12):

12) Whoever Alpie is referring to with “she” will be in the ballroom of the Palmer House Hilton at 1:05.

Betty can extract (12) because even though she doesn't know who "she" referred to in Alphonse's utterance, she knows enough about Alphonse's utterance, the surrounding facts about word meanings, and relevant features of context to know that Alphonse's utterance expressed some proposition of the form, *x will be in the ballroom of the Palmer House Hilton at 1:05pm*, and that which such proposition it expressed depends on how the context-dependence of "she" was resolved in Alphonse's context. Because she's assuming that Alphonse is speaking truly, Betty can conclude that whichever such proposition Alphonse's utterance expressed, it's true. Which gives us (12).

Betty can get to (12) on the basis of some background assumptions (about, crucially, the character of "she") that constrain, but don't uniquely determine, the proposition expressed by Alphonse's utterance. They do that because they constrain, but don't determine, the semantic value of "she" in Alphonse's context. Her partial semantic information doesn't allow her to work out the proposition expressed, but it does allow her to work out a condition that has to be satisfied in order for Alphonse's utterance to be true.

While the fact that this condition is satisfied is potentially useful information, this still looks like it's going to be a case of communicative failure. On the most natural ways of filling in the case, Alphonse probably meant to communicate the proposition about Carrie, and was expecting Betty to have the information she needed to resolve the context dependence of "she". Since Betty didn't have at her disposal all of the information that Alphonse was relying on her to have, Betty hasn't taken on all of the information that Alphonse (probably) wanted to convey.

Not every case of updating on the basis of partial information about meaning is a communicative failure, though. Here is a variation on a familiar example: There's a knock on the door. None of us can see who's there, but Bob, who's well informed about the sorts of people who are likely to be going around knocking on doors in this neighborhood at this time of day, says, "that's an encyclopedia salesperson". Since none of us can see who's at the door, neither we nor Bob are in a position to work out the proposition expressed. But we can extract a condition on Bob's speaking truly from our partial information: In order for Bob to be speaking truly, whoever is properly positioned to be the referent of Bob's use of "that" must be an encyclopedia salesperson. From this, together with some information about the particular conversational

situation that constrains the referent of ‘that’, plus the assumption that Bob is speaking truly, we can extract the information that whoever is knocking on the door is an encyclopedia salesperson.

This looks like a case of communicative success. The information that Bob was trying to convey was conveyed, because all of the background information that Bob was relying on us to bring to bear was in fact brought to bear.²⁹

This phenomenon of hearers updating on the basis of partial semantic information—information that falls short of what’s required to fully determine the proposition expressed by the utterance—is one that’s very amenable to being theorized about in an informational framework according to which all update is goat update. We’ll have occasion to talk about these sorts of updates based on incomplete semantic information a lot. So we’ll introduce a name and an acronym for them: we’ll call them *partial information updates*, or PIUs.³⁰ (What marks off PIUs as a category is the partiality of the updater’s *semantic* information, in the sense that the information’s not rich enough to fix the proposition expressed. It’s not just partiality in some respect or other. Every update by everyone except God is based on information that’s partial in some respect, so that kind of partiality doesn’t mark off an interesting category.)

The phenomenon of hearers extracting information from utterances despite lacking full information about the meanings of all the expressions used, based on an assumption that the speaker’s utterance is true (that is, that whichever proposition the utterance expressed is true) is familiar, and serves to explain many of the updates in our original example. It’s familiar, most famously, from the work of John Perry (along with a number of co-authors)³¹ and Robert Stalnaker.³² Something that we think is helpful about the goatist framework is that it offers a unifying framework that captures both Perry’s and Stalnaker’s insights about the phenomenon that (we maintain) they are both talking about. So we will pause for a moment here to briefly set up Perry’s and Stalnaker’s approaches to PIUs.

²⁹ The familiar case this is a variation on is Stalnaker’s “that’s Zsa Zsa Gabor or Elizabeth Anscombe” example in (Stalnaker, 1978), which he uses to illustrate and motivate his account of diagonalization. We’ve modified it in order to make it an instance of just one of the two phenomena Stalnaker was using his example to illustrate (that the common ground isn’t informationally rich enough to determine a unique proposition expressed) in isolation from the other (the necessary truth or necessary falsity, and consequent un informativeness, of whichever proposition is in fact expressed), because we find the squishing together of the two phenomena distracting, and we’re at present only interested in talking about the first one. (The second will come up later.)

³⁰ To be pronounced “pew”, like the sound of blaster fire in the Star Wars movies.

³¹ See for example (Korta & Perry, 2011; Perry, 1988, 2001).

³² See for example (Stalnaker, 1978, 1987, 1998, 2006, 2014)

Perry's framework is usefully thought of as beginning with a general characterization of the notion of a *truth-condition* for an utterance, such that the proposition an utterance semantically expresses in context is just a special case of a more general type. Consider again the proposition *that whoever Alphie is referring to with "she" will be in the ballroom of the Palmer House Hilton at 1:05* - (12) from our original example. This isn't the proposition expressed by Alphie's utterance, but it is something that it's plausible to call a truth-condition of Alphie's utterance, because it specifies a condition that has to be satisfied in order for Alphie's utterance to be true. It's an instance of a general type that Perry characterizes in his (very goat-friendly) gloss on truth-conditions generally: "When we ask about the **truth-conditions** of an utterance, we are asking about what else the world has to be like, for the utterance to be true, given that it occurred and meets certain conditions." (Perry, 2001), p18)

This is, importantly, a conception of truth-conditions according to which utterances have a number of distinct truth-conditions associated with them. Since a truth-condition is, generally (and rephrasing a bit in ways we think are harmless), what else has to be true in order for the sentence to express a truth, *given that the utterance met certain conditions*, we'll get different truth-conditions depending on what the "certain conditions" are that we're assuming the utterance met. For example, (12) gives a truth-condition for Alphie's utterance, because it's what else has to be the case in order for Alphie's utterance to be true, given (approximately) that there is some x such that Alphie's utterance expressed the proposition that x will be in the ballroom of the Palmer House Hilton at 1:05, and x is the person referred to by Alphie's use of 'she'.

Different choices of given conditions will determine different truth-conditions for an utterance, since what else the world has to be like in order for an utterance to be true, given P, will in general be different from what else the world has to be like in order for it to be true, given Q. We will adopt the following canonical specification of a Perry-style truth-condition for some utterance U: *What's required for U to be true, given Q, is P*. That's got three important moving parts—the utterance U, the truth-condition P, and the given information (the condition the utterance is assumed to satisfy) Q. We'll get different truth-conditions for U depending on what we build into Q. We'll sometimes talk in what follows about *the truth condition of U given Q*—what we'll mean by this is the Perry-style truth condition of U that specifies what else has to be true in order for U to express a truth, given (holding fixed) Q. Talking about the truth condition for U, given Q, is another way of talking about what we've been characterizing as the

information extractable from the utterance by bringing to bear [Q + Truth]—in general, when P is a truth condition of U given Q, a hearer will be able to infer P from the fact that the utterance occurred, plus Q, plus the assumption that U is true.

Perry, together with various co-authors, makes use of a rich taxonomy of different types of truth-conditions for utterances.³³ Not all of the distinctions incorporated into that taxonomy will be relevant for our purposes here, but we do want to pause to identify a couple of categories that will be important.

One important category of truth-conditions in Perry's taxonomy is *referential truth-conditions*, which build in the reference (more generally, the semantic values) of all of the expressions in the utterance. So the referential truth condition is just what else has to be true, given the facts about the meanings of all of the expressions in the utterance (and so, given the facts about which proposition was expressed by the utterance), in order for the utterance to be true. That is, it's just the proposition expressed. (Because P is exactly what has to be true, given that U expressed P, in order for U to be true.) Henceforth we'll sometimes use "proposition expressed" and sometimes "referential truth-condition" to talk about this. Assuming that an utterance expresses a unique proposition, it will have a unique referential truth-condition.³⁴

Another category from Perry's taxonomy that will be useful to have in our toolkit is *reflexive truth-conditions*. A reflexive truth-condition for an utterance is a proposition *about the utterance* that specifies what has to be the case in order for the utterance to be true, given that it satisfied some condition. So (12), for example, is a reflexive truth-condition for Alphonse's utterance, since it's a condition (in part) on features Alphonse's utterance has to have in order to be true—in particular, on the reference of the occurrence of "she" in the utterance. (It's important that it's *a* reflexive truth-condition—among the many truth-conditions associated with utterances in Perry's framework, a lot of them will be reflexive.)

Perry's approach to PIUs is based on attending to the very many truth-conditions associated with an utterance, differing with respect to the conditions taken as given; and observing that all of them are things that could be communicated by means of the utterance, to a hearer who is assuming that the given conditions are satisfied, and that it's true. So for example, the referential truth-condition can be communicated to a hearer who knows the semantic values

³³ See for example (Barwise & Perry, 1983; Israel & Perry, 1990; Korta & Perry, 2011; Perry, 1988, 2001)

³⁴ See for example (Bowker, 2019; Egan, 2009; von Stechow & Gillies, 2009) for pictures on which utterances may not express a unique proposition, and things will be more complicated.

in context of all of the expressions in the utterance, while a truth-condition like (12) can be communicated to a hearer who has the right kind of partial semantic information.

Stalnaker’s approach to PIUs takes as its starting point some important observations about the contingency of meaning. The truth-values of utterances are hostage to the world in two importantly different ways. For one thing, the truth-value of the proposition an utterance expresses is (typically) hostage to the world.³⁵ In addition, though, the facts about which proposition an utterance expresses are also hostage to the world. The most obvious example of this is that context-sensitivity is going to be resolved differently, depending on contingent facts about the context. But meaning is more deeply contingent than that. It’s a contingent fact that “dogs bark” means that dogs bark, rather than that pigs fly, or that fish live in water, or literally anything else. The fact that “dogs” and “bark” have the meanings that they do is contingent.³⁶

Given the contingency of meaning, what it takes for a particular utterance U to be true is that (a) the meaning-fixing facts about the world make it the case that U expresses some proposition P such that (b) the P-ish facts about the world make P true. Mia’s utterance of “Joe Biden is president” is true because (a) the meaning-fixing facts determine that Mia expressed the proposition *that Joe Biden is president*, and (b) the political facts make it true that Joe Biden is president. But Mia’s utterance could also have been true if (for example) (a) the meaning-fixing facts determined that Mia expressed the proposition that there are seven otters in Noah’s office, and (b) the otter-distribution and office-occupancy facts make it the case that there are seven otters in Noah’s office.

Generally, utterances will be associated, given the contingent meaning-fixing facts, with a particular proposition expressed. They’ll also be associated with another theoretically interesting proposition: the proposition that, whichever proposition is expressed, it’s true. This is the *diagonal proposition* associated with an utterance in Stalnaker’s framework. (So-called because it appears on the diagonal of a Stalnakerian propositional concept.³⁷) It’s the diagonal proposition

³⁵ The exceptions are cases in which the proposition expressed is either necessarily true or necessarily false.

³⁶ If you like a richer metaphysics of words, such that words have their meanings essentially, then it’s a contingent fact that those pronunciations are ways of pronouncing words with those meanings. Even more carefully: It’s a contingent fact that Bob’s production of the sounds he produced on this occasion were pronunciations of words with those meanings.

³⁷ Officially, the diagonal is only defined over worlds within the context set (Stalnaker, 2006). What we just described is really what Nico Kirk-Giannini helpfully terms the *superdiagonal* (Kirk-Giannini, 2019, 2024). But for the sake of simplicity of presentation, we’re going to go ahead and deviate from official Stalnakerian usage and talk about it as the diagonal.

(together with facts about what's presupposed in a conversation or believed by a hearer) that plays the crucial load-bearing role in Stalnaker's account of PIUs.

Because in a Stalnakerian framework, updating with the diagonal is on the explanatory ground floor of the accounts of all other PIUs,³⁸ it will be useful to take a small detour before giving the Stalnakerian gloss on updating with (12).

The diagonal proposition for Alpie's utterance appears in our initial example as (19):

19. There is some proposition P such that (a) P is the proposition expressed by Alpie's utterance, and (b) P is true.

The informationalist gloss of (19) is that it's the proposition hearers can extract from just the fact of the utterance, together with the assumption of truth. (Just from [Truth].) It's what you get from assuming just that the utterance is true, and nothing at all about the details of its meaning. In order to be true it's got to express a proposition, and the proposition it expresses has got to be true. Describing things in Perry's terms, we can think of (19) as the most-abstract truth-condition for Alpie's utterance—it specifies what the world has to be like, assuming nothing, in order for Alpie's utterance to be true.

Now we are in a position to give the Stalnakerian account of (12). The Stalnakerian will explain Betty's update with (12), in the case where her only relevant semantic ignorance is about the resolution of the context-dependence of 'she', by noting that, given what Betty is assuming about the meaning-fixing facts, (19) (i.e., the diagonal) will be true iff (12) is. Given the facts about the semantic values of the other expressions in Alpie's utterance, Alpie's utterance will express a true proposition iff Alpie's use of 'she' refers to somebody who will be in the ballroom of the Palmer House Hilton at 1:05. And so, accepting the diagonal—that Alpie's utterance expresses some truth or other—requires, given the assumptions about meaning that Betty is holding fixed, accepting (12).³⁹

According to the informationalist, what's happening with PIUs is that the hearer brings to bear some information that doesn't uniquely determine the proposition determined by the

³⁸ It's potentially a little bit misleading to call the update with the diagonal a PIU, because it doesn't require bringing to bear any semantic information at all. It's tempting to multiply pronounceable acronyms by calling it a *ZIU* (zero information update), but we will resist temptation and treat it as a limiting case of a PIU.

³⁹ Stalnaker often describes this in terms of the diagonal agreeing in truth-value with various other propositions, given what's presupposed.

utterance, but does determine a condition that has to be satisfied in order for whichever proposition it expressed to be true (for example (12)). If the hearer trusts that the speaker's utterance is true, then they can infer that this condition is satisfied.

That's also what's happening according to Perry and Stalnaker, though they describe it in different theoretical frameworks, using different terminology. Given that the audience has partial semantic information that fixes the semantic values of all the expressions in Alpie's utterance except "she"—let's call this information 'PSI'—and that they assume that Alpie's utterance is true, they will be in a position to infer (12). The informationalist describes this in terms of (12) being inferrable from U on the basis of [PSI + truth]. The follower of Perry will describe it in terms of (12) being a truth-condition of Alpie's utterance, given PSI. As such, it will be inferrable from the occurrence of Alpie's utterance, on the basis of PSI, together with the assumption that Alpie's utterance is true. The Stalnakerian will describe it in terms of (12) being a consequence of the diagonal for Alpie's utterance plus PSI, and so inferrable by an audience who observes Alpie's utterance, assumes it's true, and believes PSI.

Let's look at a couple more of the PIUs that are liable to result from Alpie's utterance:

18. There are a proposition P, person x, time t, and location l such that: (i) P is the proposition that x will be at l at t; (ii) x is the person referred to by Alpie's use of 'she'; (iii) l is the location referred to by Alpie's use of 'here'; (iv) t is a time 10 minutes later than the time referred to by Alpie's use of 'now'; (v) Alpie's utterance expressed P, and (vi) P is true.

Compared to (19), this is the somewhat more specific proposition that you get by holding fixed that the utterance is true, and also holding fixed the context-independent facts about the meanings of the expressions.⁴⁰ It's a proposition updating with which requires Betty to bring to bear information about the types of semantic values that context will determine for the context-sensitive expressions in the sentence Alpie produced,⁴¹ but doesn't require her to bring to bear any substantive information about how that context-sensitivity is resolved. It's extractable

⁴⁰ Maybe not quite. (18) doesn't incorporate facts about the specific pattern of dependence of semantic value on content that will likely also be part of context-independent meaning—which is encoded in an expression's Kaplanian (Kaplan, 1989) character, for example.

⁴¹ That is, information about whether the semantic values will be times, locations, etc.—not information about whether they'll be of semantic type $\langle e, t \rangle$, $\langle \langle e, t \rangle, t \rangle$, etc.

from Alpie's utterance on the basis of [context-insensitive semantic values + types of context-sensitive semantic values + truth]. It's also a truth-condition of Alpie's utterance—it's what else has to be true, given the semantic contributions of the context-independent expressions, and the types of semantic values the context will fix for the context-dependent ones, in order for Alpie's utterance to be true. And it's a proposition that, given the semantic contributions of the context-independent expressions, and the types of semantic values the context will fix for the context-dependent ones, will be true if the diagonal is.

A still more specific truth-condition is (17) - here what we're assuming about the utterance imposes all the same constraints as gave us (18), and also fixes the reference of "now":

17. There are a proposition P, person x, and location l such that: (i) P is the proposition that x will be at l at 1:05pm; (ii) x is the person referred to by Alpie's use of 'she'; (iii) l is the location referred to by Alpie's use of "here"; (iv) Alpie's utterance expressed P, and (v) P is true.

That's also a proposition, about Alpie's utterance, that specifies how the world has to be for it to be true, given that it occurred and meets certain conditions. It differs from (18) by building more into the "certain conditions" that we're assuming Alpie's utterance meets, such that we're asking what more beyond that is required in order for Alpie's utterance to be true. It's a proposition extractable from Alpie's utterance on the basis of the information that supported (18), plus the facts about the resolution of the context-dependence of 'now'. It's what else has to be true in order for Alpie's utterance to be true, given the same background conditions that gave us (18) as a truth-condition, plus the resolution of the context-dependence of 'now'. And it's a proposition such that, given the information that supported the local entailment of (18) by the diagonal of Alpie's utterance, plus the resolution of the context dependence of 'now', it will be true if the diagonal is.

A running theme in Perry's work, along with his co-authors', is the importance of truth-conditions other than referential ones in accounting for how we learn from, and communicate with, utterances. This is, we think, an extremely important insight, and Perry and co. make use of it to do a lot of work in explaining interesting communicative phenomena, dissolving puzzles, and warding off potentially tempting mistakes. We would like to, and we

think that we can, endorse pretty much all of this, because we think that Perry's fundamental insight about the multiplicity of truth-conditions associated with utterances, and their communicative importance, is extremely goat-friendly. The goatist gives an account of updating with the fact that a particular utterance occurred, while also bringing to bear an assumption that the utterance is true, plus some additional information that may or may not determine the proposition expressed by the utterance. Perry gives an account of updating on the basis of some information about conditions an utterance satisfies, which may or may not determine the proposition expressed by the utterance; knowledge of what else is required, given that the utterance satisfies those conditions, in order for it to be true; and the assumption that it's true. These look to us like two modes of presentation of the same phenomenon.

A running theme in Stalnaker's work is the importance of updates, not with the proposition in fact expressed by an utterance, but with the proposition that the utterance expressed some truth or other, in accounting for how we learn from, and communicate with, utterances. This is, we think, an extremely important insight, and Stalnaker makes use of it to do a lot of work in explaining interesting communicative phenomena, dissolving puzzles, and warding off potentially tempting mistakes. We would like to, and we think we can, endorse pretty much all of this, because we think that Stalnaker's fundamental insight about the communicative force and value of the information that the world satisfies the conditions it needs to satisfy in order to make it the case that, whichever proposition an utterance expresses, it's true, is extremely goat-friendly. The goatist gives an account of updating with the fact that a particular utterance occurred, while also bringing to bear an assumption that the utterance is true, plus some additional information that may or may not determine the proposition expressed by the utterance. Stalnaker gives an account of updating with the proposition that a particular utterance expressed a truth, given a belief-state (or a context-set) including some other information that may or may not determine which proposition the utterance expressed. These look, again, like two modes of presentation of the same phenomenon.

We think that the fact that Stalnaker and Perry describe the phenomena that they're theorizing about in such different terminology, working in such different theoretical frameworks, obscures the fact that they're theorizing about the same phenomenon, and often making very similar claims about it. Something that we have found helpful about framing our theorizing about PIUs in informational terms is that it makes it easy to see that they *are* theorizing about the

same phenomenon, and that at least the fundamental ideas behind their approaches to it are compatible with each other (and, we think, both importantly correct). What both Perry and Stalnaker are talking about is the phenomenon of PIUs based on an assumption of truth. And Perry's right that it's extremely valuable, when thinking about what we can learn from (and communicate with) an utterance, to think about what audiences know (or believe) about which conditions the utterance satisfies, and what else has to be true, given that the utterance satisfies those conditions, in order for the utterance to be true. And Stalnaker is right that it's extremely valuable, when thinking about what we can learn from (and communicate with) an utterance, to pay careful attention to the role of the world in determining which proposition the utterance expressed, as well as the role of the world in determining the truth-value of the proposition that was in fact expressed.

So we think the informationalist approach is helpful for making it easy to see that Stalnaker and Perry are theorizing about the same thing (PIUs based on an assumption of truth), and that at least the fundamental ideas behind their approaches aren't in tension. (Though maybe there are some specific applications that are.) It's also helpful in making it clear that we shouldn't be at all surprised to see the sorts of phenomena Perry and Stalnaker are theorizing about, and that the phenomenon is separable from the particular theoretical tools that Perry and Stalnaker use to characterize it. Of course we can make inferences based on the occurrence of an utterance, plus partial semantic information, plus the assumption that the utterance is true. And what we'll infer then is exactly the sort of thing Perry and Stalnaker talk about—we'll infer that things are the way they need to be, given the partial semantic information available to us, in order for the utterance to express a truth. (So we shouldn't, for example, be suspicious of PIUs as a communicative phenomenon because we're suspicious of propositional concepts, or because we have objections to Perry's taxonomy of truth-conditions.)

We also think that the informationalist account is helpful in two more ways: First, it makes it easy to notice and draw attention to some features of the sorts of PIUs that Perry and Stalnaker drew attention to that are important, but are perhaps not as obvious under Perry's or Stalnaker's modes of presentation. And second, it makes it easy to notice and draw attention to a range of PIUs that Perry and Stalnaker don't talk about, but which are also theoretically interesting and communicatively important.

4.5 Some further morals about PIUs based on truth

The first item on our agenda in this section will be to emphasize, illustrate, and support some points about PIUs, the communicative intentions of speakers, and what makes for communicative success and failure, that have already come up in passing.

Another pair of points that we'll be concerned to reinforce throughout this section are that (a) PIUs are ubiquitous, and (b) we shouldn't be at all surprised, if we're working in an informationalist framework, that they're ubiquitous. We think these two facts are interesting and important for at least two reasons: First, since PIUs really are ubiquitous, the fact that an informationalist approach makes their ubiquity completely unsurprising is a point in favor of informationalism. And second, given their unsurprising ubiquity, we shouldn't regard appeal to PIUs in explanations of the communicative force of utterances as any kind of suspicious epicycling or special pleading.

We begin with a couple of examples from Perry, plus one from Grice (the first Perry example is verbatim, the other two are slightly modified versions of the originals):

“Ellsworth goes to Hawaii and sends me a postcard. Unfortunately, it gets a bit wet before I receive it. The postmark, return address, and signature are all illegible. The message stays dry: ‘I am having a good time now’.” (Perry, 1988 pp)

An attentive driver, whose eyes are locked to the road, hears shouts of “stop there” and “that’s a good place” from the back seat of the car, but is unable to perceive the speaker’s accompanying demonstrations. (Perry, 1988 pp)

After getting distracted and losing track of what was happening in the conversation for a while, you snap back to attention to hear Paul saying, “He is in the grip of a vice”. (Grice, 1975)

A point that Perry and Grice are concerned to make with all of these examples is the by now familiar one that there’s some information (distinct from the proposition expressed) which one can extract from utterances even if one doesn’t have enough information to work out which

proposition the utterance expressed. In the smudged postcard case, the reader can learn, based on an assumption of truth, that whoever wrote the postcard was having a good time at the time when they wrote it. In the backseat demonstrative case, an extremely relevant PIU is that there's a place within demonstrating-distance that the speaker wants to stop at. In Grice's case, it's that there is "some particular male person or animal x [who Paul was intending to talk about], that at the time of the utterance... either (1) x was unable to rid himself of a certain kind of bad character trait or (2) some part of x's person was caught in a certain kind of tool or instrument". (We also repeat Grice's caution: "approximate account, of course".)⁴²

It should be clear that this is all extremely goat-friendly, and the informational account of these cases is straightforward. The hearer lacks some of the information that they need in order to zero in on the proposition expressed, because they lack certain information necessary to execute the compositional process that determines the expressed proposition. One reason why a hearer might not be able to complete the compositional process is because they don't know enough to resolve the context-dependence of one of the expressions used, and so they don't know which semantic value to plug in at some point (or points) in the compositional process. (Which object of type $\langle e, \langle s, t \rangle \rangle$ does "tall" or "flat" or "ready" contribute in this context? Which object of type e does "he" or "she" or "that guy" contribute?) That's what's happening in the two examples from Perry. The hearer has enough information to extract a character-fixed truth condition (e.g., there's a place x such that x is visible from the car and x is a good place to stop), but not enough to extract a referential truth condition (e.g., that Reuben's Diner is a good place to stop).

Many of the examples in the literature, and much of our initial discussion of Alphonse's utterance, are focused on PIUs based on semantic information that's partial because it's missing information about the resolution of the context-dependence of some expression(s). It's worth emphasizing, and Grice's example illustrates, that ignorance about the resolution of context dependence is just one way in which one can fail to be in a position to complete the compositional derivation of the proposition expressed. Another reason is that one might not know which sentence was produced. This might be because of a lexical ambiguity—for some part of the utterance, you know the sound or the written shape tokened, but you don't know

⁴² We deviate from Grice's intended use of the "in the grip of a vice" example by focusing on what hearers can learn based on assuming that Paul spoke truly, whereas Grice focuses on the interpreter's partial information about what the speaker said.

which lexical item it was a tokening of. Or it might be because you know which lexical item was produced but you don't know (or don't know exactly) what it means. Or it might be because one of the words on the postcard was smudged and illegible, or because some part of the speaker's utterance was inaudible due to the tolling of the church bells, or a passing train, or a zoom glitch, or a nearby partygoer's coughing fit.

These examples from Perry and Grice all look like cases of at least partial communicative failure—they're all cases where the speaker had a particular communicative aim, and the hearer didn't have enough information to pick up on a proposition that the speaker was intending to communicate. Many PIUs are like this, but many are not—many look like straightforward cases of communicative success.

The encyclopedia salesperson case discussed earlier is an example, as is the “that's Zsa Zsa Gabor or Elizabeth Anscombe” case from Stalnaker (1978) that it was based on. These are cases in which what the audience is in a position to learn from the utterance, on the basis of their partial semantic information and the assumption of truth, is a mix of information about meaning and information about non-semantic features of the world. The truth of the utterance imposes constraints on the reference of “that”—it must be an encyclopedia salesperson, or one of Gabor or Anscombe. Given some additional information that further constrains the reference of the use of ‘that’ in the utterance—that it refers to whoever just knocked on the door, whoever just passed by in the hallway, etc.—the audience can extract a constraint on entirely non-semantic features of reality: That whoever just knocked on the door is an encyclopedia salesperson, or that the person who just passed by in the hallway is either Gabor or Anscombe. In these cases (at least on natural ways of filling them in), the speaker's communicative agenda is entirely concerned with this non-semantic information, and they're cases of communicative success because the information that the speaker intended to convey was in fact received.

In another important class of cases of PIUs as successful communication, it's the semantic information that's the target of the speaker's communicative intention. Another famous case from Stalnaker (1978) serves as an example. Andy doesn't know what ‘ophthalmologist’ means, and so Tim tells him, “an ophthalmologist is an eye doctor”. Because Andy doesn't know what ‘ophthalmologist’ means, he can't extract the proposition in fact expressed by Tim's utterance. But that's fine—based on the assumption that Tim's utterance is true, Andy can learn that ‘ophthalmologist’ has a meaning such that ‘an ophthalmologist is an eye doctor’ expresses a

truth. Given Andy's knowledge of the meanings of the rest of the words, this is likely to be enough to allow Andy to learn, on the basis of Tim's utterance, what 'ophthalmologist' means,⁴³ which was exactly what Tim intended to convey.

One reason why this kind of case is important is that it looks like a straightforward case of fully successful communication. Another is because it highlights the fact that many of the potential updates associated with utterances will be updates with *linguistic* information—information about linguistic features of the utterance, and/or of the expressions that occur in it. This should be completely unsurprising on an informational account—assuming that an utterance is true, while holding fixed the meanings (and syntactic roles) of some, but not all, of the expressions that occur in the utterance, will systematically impose constraints on the meanings (and the syntactic roles) of the remaining expressions.

Examples of this kind of update are not hard to find. If Mathew hears Mandy say, "I [unfamiliar expression] to the seminar with Craige yesterday", he can work out that the unfamiliar expression must be a verb, and maybe that it's likely to be a verb having something to do with modes of transportation. (Note, by the way, a fact that we'll return to a little ways downstream—that Matthew doesn't actually need to be assuming truth to get this information. Grammaticality and relevance will do.) We also shouldn't be surprised that speakers are liable to notice and exploit the fact that such information can be extracted by hearers.

Exactly the same mechanism is at work in the interpretation of partially-understood utterances in languages the interpreter is not fully fluent in: As a visitor in Germany, given Andy's extremely partial knowledge of German, it's very likely that he will encounter some utterances that contain a mix of familiar and unfamiliar vocabulary. (We introduce an example containing some non-words in order for it to work as a partially-understood sentence even for readers whose German is much better than ours. Please join us in pretending that the words in all caps below are actual German words.)

Suppose a helpful local guide tells me, "hier muss man der SCHNITZELSTANGER VERGLUPTZEN". If I don't know "SCHNITZELSTANGER" or "VERGLUPTZEN", but I do know "hier muss man," and a bit about what German nouns and verbs tend to look and sound like, what I can extract is *that there's something one's gotta do something to in that location*. (If I

⁴³(Stalnaker, 1978). Note that confidence that Tim is speaking truly doesn't quite guarantee that Andy can work out exactly what 'ophthalmologist' means - 'an ophthalmologist is an eye doctor' could also be true if 'ophthalmologist' referred to a specific type of eye doctor. (Thanks to Wolfgang Schwartz for pointing this out.)

know enough to resolve the context-dependence of “hier”, it’ll be *there’s something one’s gotta do something to at L*. If not, it’ll be *there’s something one’s gotta do something to in the F location*, where F is the strongest constraint I can work out on the resolution of “hier”.) Being able to do this kind of thing is sometimes super helpful for figuring out the meanings of unfamiliar vocabulary. It’s also sometimes helpful for narrowing down the range of things to do if one wants to act in accordance with local norms and expectations. These kinds of cases of coping with one’s partial knowledge of a second language are likely to be cases in which the hearer fails to perform the update the speaker was intending to produce, and so they’ll be cases of hearers trying to make the best of what they recognize as a case of communicative failure.

Other very similar cases, though, are cases of communicative success. A lot of immersion-style language instruction happens this way. The teacher leverages the parts of the target language that the students have already learned in order to produce sentences containing parts they don’t yet know, in a context in which the students can work out what the unfamiliar expressions must mean in order for the teacher to be saying something true.

The learner is also very likely—to preview a point we’ll be emphasizing later—to rely on many other assumptions besides truth in extracting information about the meanings of the novel expressions. For example, assuming that the meaning is such as to make what the teacher is saying relevant, or appropriately related to the pictures he’s pointing to, or such that the novel expressions are the sort of thing it would make sense to teach at this point in the lesson, is also likely to be helpful in narrowing down the range of potential meanings. (It’s also probably worth noting that it’s not just new vocabulary, but also things like verb conjugations, word order rules, tenses, etc. that get taught this way.)

The same thing happens with unfamiliar technical terms. One of the contributors to the impostor syndrome that plagues many of us in our first year(s) of graduate school is this kind of failure to fully understand what people are talking about because one doesn’t know the meanings of all the words they’re using. One way out of this is to explicitly ask a lot of questions about what all this strange new terminology means. But it’s often useful to be able to work out some constraints on what the users of the unfamiliar vocabulary must be saying (for example when it’s the middle of a talk and it’s not time for questions yet, and the speaker skipped, or you were distracted during, the part where they introduced their technical terms). And that’s something we can do. We do it by exploiting our knowledge of what the other words in the utterance mean,

whatever partial knowledge we might have of the meaning of the un-mastered technical term, and whatever information we have about the surrounding context in which the utterance took place. This ability to leverage information from facts about the utterance plus background knowledge, together with an assumption of truth, is an important part of a healthy response to the first year graduate student predicament (or the attending-a-talk-in-an-unfamiliar-subfield predicament, or quite generally the trying-to-learn-something-new predicament): you can learn a lot of technical vocabulary by working out what it would have to mean, given what else you know, in order for the speaker's or writer's contribution at that point in the conversation or the paper to be true. (More generally: in order for the contribution to have some feature you trust that it has—perhaps truth, but also potentially sincerity, QUD-responsiveness, combination with some other utterances to produce a valid argument, being such that if it were true the objection under discussion would be defused, etc. Here again, skilled application of this strategy involves holding fixed a lot of other features of the utterance in addition to—and sometimes instead of—truth..)

Some morals about communicative success and failure: The difference between cases of communicative success and and communicative failure is whether the audience updates in the way the speaker intended. Whether this happens or not will typically depend on whether the audience brought to bear the information that the speaker was relying on them to bring to bear. Thinking about communication in an informational framework, there's no reason to expect that the audience's updating with the proposition semantically expressed by an utterance should always be among the speaker's communicative aims. So it should not be at all surprising that we see utterances such that the audience doesn't update with the proposition expressed, but which are nonetheless fully communicatively successful. Nor should we be surprised to see utterances such that the audience *does* update with the proposition expressed, which are nonetheless communicative failures. (Cases of un-picked-up-on implicatures are obvious examples.) And given that there will often be more than one update that a speaker is aiming to produce with a single utterance, we should expect to see varying kinds and degrees of partial communicative success and failure, in which the audience picks up on some, but not all, of the updates that the speaker was aiming to produce.

Another class of cases worth drawing attention to, which features prominently in discussions by both Stalnaker and Perry, are ones in which the communicative force of the

utterance has both an important linguistic component and also an important component that's about non-linguistic features of the world. Both Stalnaker's and Perry's responses to Frege's puzzle cases have this structure. Take the case in which John says to Bob, "Hesperus is Phosphorus".

Here is a piece of information that Bob can extract from John's utterance, given the assumption that John is speaking truly, and some basic information about the kinds of expressions 'Hesperus', 'Phosphorus', and 'is' are: *The thing that satisfies the condition of reference for John's use of "Hesperus" is the thing that satisfies the condition of reference for John's use of "Phosphorus"*.

That's a proposition that follows from the diagonal given the additional information that "Hesperus" and "Phosphorus" are names and "is" (or the occurrence of "is" in the target utterance) means identity. It's also the truth-condition of John's utterance, given that same additional information.

If Bob has more specific beliefs about what the conditions for reference of the expressions are (for example, that "Hesperus" refers to whichever object is appropriately causally and historically connected to John's use of "Hesperus"), he'll be able to get some more specific information. One way to represent this is in terms of the difference between the truth-condition for John's utterance mentioned in the previous paragraph, and the one that also includes in its given conditions a substantive characterization of the reference relation. Another way to represent it is in terms of the difference in what can be inferred from the diagonal plus the information above about "Hesperus", "Phosphorus", and "is", and what can be inferred from the diagonal together with that information, plus a substantive characterization of the reference relation.

He will also be able to extract more information if he has some additional beliefs about which other properties the reference-condition-satisfiers have. For example, if he thinks that the object that satisfies the reference-condition for "Hesperus" is the first celestial object visible in the evening, and the object that satisfies the reference-condition for "Phosphorus" is the last celestial object visible at night. And again, we'll be able to describe this either in terms of different Perry-style truth conditions, or in terms of differences in what follows from the diagonal, together with different packages of background information.

The informationalist will follow Perry's and Stalnaker's recipe for treating Frege's puzzle cases as cases of successful communication of substantive, contingent information, even though the proposition expressed is necessarily true, because they'll follow Perry and Stalnaker in treating them as cases in which the speaker's communicative intention is targeting one or more PIUs associated with the utterance, rather than the proposition semantically expressed. (Though Perry and Stalnaker might not describe their proposals in exactly that way.) The informationalist will follow both Perry and Stalnaker here because, for the informationalist, it will look like Perry and Stalnaker are giving different descriptions of the same (correct!) response to Frege's puzzle.

To introduce the next point that we want to make about PIUs, we return to Perry's example of the passenger in the back seat calling out, "that's a good place to stop!". In Perry's original presentation, this is clearly a communicative failure. The speaker is trying to communicate information about a specific restaurant, which the driver can't extract on the basis of the information he has access to. The PIU the driver performs is a sort of second-best mechanism for coping with the failure of communication. Two things to note about this case:

First, we can easily imagine that in a case like this, although the speaker is in fact attempting to communicate information about a specific restaurant, securing this kind of second-best update with a reflexive truth-condition is enough to secure success in the practical project that driver and passenger are pursuing. If all the restaurants within gestural-demonstration distance are off the same highway exit, then knowing that there's some restaurant the speaker is demonstrating that they want to stop at is enough to resolve the question of what the driver ought to do right now: he should take the exit. (Some additional, more fine-grained communication will then probably need to happen in order to get all the way to the restaurant.) The driver's interpretive damage-control doesn't secure fully successful communication, but it does allow for the same kind of practical coordination that fully successful communication would have enabled.

This points the way to a second thing worth noticing: we can *also* easily imagine a version where eliciting the right PIU is precisely what the speaker was aiming for. In other words, where the driver's updating with the PIU doesn't constitute damage-control, but rather straightforward communicative success. This is the case where all the speaker was trying to convey to the driver was the information that there's some restaurant within gesturing range that it'd be good to stop at (and so he should take the exit). In this case the driver still can't extract

the proposition semantically expressed, but that's no barrier to communicative success, because the speaker wasn't attempting to communicate the proposition semantically expressed, just the PIU (which, the speaker correctly surmised, was enough to properly guide the decision—take the exit or not—that was immediately at issue).⁴⁴

And this in turn points to the main observation we are concerned to make with this example: PIUs are often communicatively important, and represent significant additional information, even when the hearer *is* in a position to update with the referential truth condition. (Even when they're in a position to perform what it's natural to call a *full information update*, or FIU.⁴⁵)

Consider the case in which the passenger is in the front seat, and the driver can observe the demonstration without irresponsibly diverting their attention from the road. The driver will then be in a position to identify the particular restaurant that's the referent of “that” in the passenger's utterance of “that's a good place to stop”. He will therefore be in a position to update with the utterance's referential truth condition—for example, *that Pete's Bar and Grille is a good place to stop*.

Notice that this FIU, by itself, is *not* enough to motivate taking the exit if the driver doesn't already know where Pete's is. The passenger could have produced the same update, four exits earlier, by saying “Pete's Bar and Grille is a good place to stop”, and that wouldn't have made it rational for the driver to take any particular exit, absent further information about the location of Pete's Bar and Grille. Just giving the information about Pete's being a good place to stop doesn't rationalize taking any particular exit in the absence of some information about where Pete's is. Given that (a) information about Pete's location is no part of the proposition semantically expressed by the passenger's utterance, and (b) we may assume that the driver didn't already know where Pete's is, it must be that if the utterance makes it rational to take the exit, it will be because it succeeds in communicating to the driver, not just the information that Pete's is a good place to stop, but also some information about Pete's location.⁴⁶ And the way the

⁴⁴ Some other, parallel cases, in which it's likely not to be important for the speaker's communicative purpose that the hearer extract the specific proposition expressed: “he's got a gun!”, “that's a lion!”, “the exam isn't today, it's tomorrow”. See also the cases in (Buchanan & Ostertag, 2005; King, 2021)

⁴⁵ To be pronounced “phew,” like the sound you make when you realize you have all, or even more than, the information you require.

⁴⁶ We are inclined (again one more emphatically than the other, but this time in the other order from in footnote 7) to insist that it has to communicate to the driver some essentially self-locating information, which they might describe as “information about where Pete's is relative to me”. But that's not essential for the current point. For present

utterance is going to communicate that kind of locational information is going to be by way of a PIU (for example, that the restaurant at location *l* is a good place to stop, where *l* is the location the passenger is pointing at).

That's a case where, very plausibly, the FIU and the PIU are both non-redundantly communicatively helpful. As we've seen, even when the hearer is in a position to extract the FIU, the PIU is still an independently valuable contribution. Indeed the PIU might (as mentioned in our initial discussion of the example) be enough to get the driver to take the exit. But the FIU is also likely to be independently helpful—for example if they get lost after taking the exit and need to ask directions, or want to make recommendations to their friends later about good places to stop. And it's certainly not entailed by the PIU, so it also constitutes a non-redundant informational contribution.

This illustrates what we think is an important point, that FIUs and PIUs need not be in competition with one another as potential communicative effects of utterances, or as objects of speakers' communicative intentions. It's perfectly possible for hearers to update with both, it's perfectly possible for both updates to be independently communicatively valuable, and it's perfectly possible for speakers to intend both. There are also no uniform relations of communicative priority between the two. It's not, for example, the case that whenever the FIU is available, producing it must be the speaker's primary communicative aim. It's not even true that, whenever the FIU is available, producing it must be among the speaker's communicative aims.⁴⁷

To drive home this last point, consider that there are also many cases in which a PIU is going to be communicatively helpful, while the FIU, even if it's available, is not. Consider Perry's (2001) case of David Lewis coming to the train station to pick up a visitor. Lewis approaches the visitor and says, "I am David Lewis". The proposition in fact expressed—that David Lewis is David Lewis—is not especially communicatively helpful. But there are some PIUs associated with Lewis's utterance that are extremely communicatively helpful, for example that the person who just spoke is David Lewis. The communicative value of introductions is entirely bound up with the PIUs associated with them—with information that one doesn't require

purposes, it's enough that in order to rationalize taking the exit, the utterance has to communicate to the driver some information about, for example, where Pete's is relative to exit 23.

⁴⁷ Here we pick a bit of a fight with Stalnaker, or at least with a natural reading of Stalnaker, on which diagonalization is treated as always a repair strategy, to be broken out only in cases where the proposition in fact expressed is underdetermined by the common ground, or fails to distinguish between possibilities compatible with the common ground. See Stalnaker 1978, and also the exchange in (Almotahari & Glick, 2011; Hawthorne & Magidor, 2009; Stalnaker, 2009)

full semantic information to extract from them. This is true regardless of whether the hearer has, in addition to the information required to extract the relevant PIU(s), the information required to determine the proposition expressed. Even in cases where the visitor can work out the proposition expressed,⁴⁸ that won't be the information that is communicatively helpful.

The introduction cases also have nearby relatives in which both the FIU and a PIU are intended, and neither is redundant given the other. David Lewis, expecting to be recognized, tells the visitor, "I am here to drive you to the department". Both the information that David Lewis is at the train station to drive Vanessa to the department, and the information that the author of *u* is at the location of *u* to drive the addressee of *u* to the department are likely to be independently useful pieces of information for the visitor to have. (For example, the PIU that the speaker is here to drive the addressee to the department is going to be immediately helpful in determining that now is the time to pick up her bag and follow the speaker to his car, while the FIU is likely to be helpful in deciding what areas of philosophy to talk about during the drive.) Lewis should expect both to be extractable on the basis of his utterance and his addressee's information, and he may well intend to communicate both.

We'll close this section by recapitulating what we take to be the most helpful morals about PIUs that approaching from a goatist angle makes it easier to see:

First: The phenomenon is totally ubiquitous. Updates on the basis of partial semantic information happen *all the time*.

Second: The phenomenon is not at all surprising, and it's not at all surprising that it's ubiquitous. *Of course* these kinds of updates happen all the time. For one thing, whenever we're in a position to update on the basis of full semantic information, we'll also be in a position to update on the basis of partial information. For another thing, it happens all the time that the best we can do is partial information. We are very often faced with utterances in which people use unfamiliar words, we miss part of the sentence, or we're missing some information about context relevant to resolving context-sensitivity. These cases are not at all exotic. We all find ourselves in this kind of situation all the time, and when we're in those situations we perform PIUs. In addition, there will systematically be very many PIUs associated with a single utterance, given

⁴⁸ If they pick up on the target PIU, they'll definitely be able to determine the proposition expressed after, but doing so won't give them any useful additional information. And they might be able to determine the proposition expressed without going by way of the PIU if, for example, they recognize Lewis's voice.

the information available to a hearer, corresponding to Perry-style truth-conditions that take different subsets of the total information available to the hearer for granted.

Third: The very many PIUs associated with an utterance are not in competition with each other, and are not in competition with updates based on full semantic information. Relatedly, PIUs are not in general washed out by or redundant in the presence of updates based on full semantic information. It's perfectly possible—in fact we should expect it to be extremely common—for speakers to intend to communicate several different PIUs with a single utterance, or to intend an update based on full semantic information along with one or more PIUs. (This is a point Perry and co-authors make in a number of places, using different terminology. See for example (De Ponte et al., 2023; Korta & Perry, 2011; Perry, 2001).)

Fourth: The underlying phenomenon, which is definitely real and definitely ubiquitous, is separable from the particular bits of theoretical apparatus people have used to talk about it. So, for example, being unhappy about propositional concepts and the dagger operator is not a reason to be unhappy about the communicative importance of PIUs.

In our experience, explicitly approaching questions about communication in informational terms makes it easier to think clearly about the phenomenon of PIUs, and easier to see how prior discussions of the phenomenon relate to each other. Another crucial benefit of approaching PIUs from a goatist angle: It makes it easy to see that truth is a bit of a sideshow. Perry and Stalnaker's approaches to PIUs have a lot in common. One thing they have in common is that they both build in the assumption of truth as a constant element of the analysis. But truth is just one of many features that we can assume an utterance has, and that we can thereby extract information from the utterance on the basis of, even when we lack complete semantic information. It's also just one of many features speakers can rely on us to assume an utterance has, and can thereby exploit in order to communicate stuff to us, even when they know that we lack complete semantic information. We develop this thought in the next section.

4.6 PIUs based on assumptions other than truth

We saw early in our discussion that we can get a lot of information and potentially valuable updates from full semantic information—information that fixes the proposition expressed—without an accompanying assumption of truth. For example, in the case of Tim's

utterance of “It’s snowing in Lexington,” we can get that it’s snowing in Lexington from [utterance + proposition expressed + truth]. But if the assumption of truth isn’t available, we can still get that Tim *believes* it’s snowing in Lexington from [utterance + proposition expressed + sincerity]. And we saw that we get something similar in the case of implicatures. When we can’t, for example, extract the canonical relevance implicature because we’re not assuming that the speaker is *succeeding* in being relevant, we can still get some interesting and potentially useful updates from the assumption that they’re *attempting* to be relevant.

We will, not surprisingly, get precisely parallel phenomena for PIUs. In the previous section, we’ve been following Stalnaker and Perry in focusing entirely on the updates that follow from some information that determines a condition that has to be satisfied in order for the utterance to be true, and also assuming that it’s true. But these truth-based updates are just a small slice of all the updates likely to be available. If we assume the speaker takes themselves to be speaking truly, or that they have good evidence for taking themselves to be speaking truly, or that they intend us to believe that they’re speaking truly, etc., we’ll get another bunch of updates, about the speaker’s beliefs, evidential situation, and/or intentions with respect to various PIUs. (Assuming also that we think the speaker associates the same PIUs with the utterance that we do.)

Consider (13) - (15) from our initial example:

13. Carrie will be in the Palmer House Hilton ballroom ten minutes after the time of Alphonse’s utterance.
14. Alphonse believes that Carrie will be in the Palmer House Hilton ballroom ten minutes after the time of Alphonse’s utterance.
15. The fact that Carrie will be in the Palmer House Hilton ballroom ten minutes after the time of Alphonse’s utterance is relevant to some goal or project that Alphonse and Betty are pursuing.

Update with (13) is supported by a package of partial semantic information strong enough to determine that Alphonse’s utterance is true iff (13) is, plus an assumption that Alphonse is speaking truly. In other words, [utterance + most of the semantically relevant information, minus what time it is + truth]. (This could be what Alphonse intends to communicate if, for example, he doesn’t know the time, and just got a message from Carrie saying, “there in 10 minutes”.)

Update with (14) is supported by that same partial semantic information, but combined with an assumption of sincerity, rather than truth.

And update with (15) is supported by the assumption that (13) is something Alphie is *intending* to communicate, which is in turn likely be supported, in part, by the same partial semantic information, along with the assumption that Alphie is being relevant. In other words, [(13) + communicative intention towards (13) + relevance]. (We take it that the *communicative intention towards 13* assumption is important here in order to license the implicature-like inference from (13) to (15).)

Note that (15) illustrates two distinct and important points:

First, it's an example of an update based, not on something weaker than truth, but on something that's just orthogonal to truth. An assumption of sincerity is naturally thought of as a weakening of an assumption of truth, because sincerity is often a part of the package of assumptions about an utterance that jointly entail its truth—for example, sincerity plus competence.⁴⁹ Relevance isn't like that.

Second, it illustrates that on an informationalist account, where very many different updates associated with a sentence are available to be targeted with a speaker's communicative intentions, we should expect to see the full range of implicatures based on assumptions of various species of cooperativeness with respect to whichever updates are (taken to be) objects of the speaker's communicative intentions.⁵⁰ So in particular we should very much *not* expect to see implicatures based only on the proposition expressed. We should expect also to see implicatures based on PIUs that are taken to be deliberately communicated.

That's good, because we pretty clearly *do* see implicatures based on PIUs. It's completely reasonable for hearers in the examples discussed in the previous subsection to expect that the information, variously, about the coreference of "Hesperus" and "Phosphorus"; about the coinstantiation of *being the first celestial object visible in the evening* and *being the last celestial object visible in the morning*; about the meaning of "ophthalmologist"; that the person who just knocked on the door is an encyclopedia salesperson, or is either Elizabeth Anscombe or Zsa Zsa

⁴⁹ On the strong, correctness-entailing understanding of competence common in some threads of the literature on testimony.

⁵⁰ Whether these are implicatures proper, or "implicature-like inferences," depends on whether implicatures are generated, by definition, from the proposition literally expressed. We take this to be mostly a terminological question, but our sympathies are with the idea that the much more natural category is the broader one: implicatures as inferences from whatever information is the target of the speaker's communicative intention, plus cooperativity.

Gabor, is relevant, quantitatively appropriate, expressed in a manner suited to the present conversational purposes, etc., and to draw further conclusions on that basis. And it's perfectly reasonable for speakers to expect hearers to do this, to exploit this fact to communicate additional messages to their audiences, and to do this in a way that's transparent in the way required for those messages to count as objects of Gricean meaning_{NN}. (It's also possible for them to satisfy less demanding conditions while still being deliberately, and even transparently, communicated.) It is a virtue of the informationalist account that it makes it easy to see that we should not be at all surprised to see this happen, and that it makes it easy to characterize and theorize about the conditions that need to be in place for it to happen.

These are situations where the Gricean mechanisms work in the usual way, but where the *input* to those mechanisms is a PIU, rather than the proposition expressed. The implicatures the hearer calculates in these cases are precisely what the speaker intended to communicate. But things can also work, in a certain sense, the other way around: The making-the-best-of-communicative-failure phenomenon Perry and Stalnaker both talk about alongside their cases of communicative success—extracting useful information from utterances in situations where it may not be possible to calculate the intended informational content—can occur with implicature, just as it does with the linguistically encoded content of the utterance.

Suppose Perry's smudged note is not a postcard, but rather a letter of recommendation. It reads, "To whom it may concern... [illegible] command of English is excellent, and attendance at tutorials has been regular. Sincerely, Daniels." It seems clear that we can get what it's natural to call "diagonal implicatures" in situations like this, where we don't know enough to work out the intended implicatum. In the case of the partially understood reference letter, with smudged subject name: There's somebody who Daniels thinks isn't good at philosophy. In the variant with smudged sender name: One of O'Leary's recommenders thinks O'Leary is not good at philosophy.

Here is what we take to be a helpful lesson of assimilating Grice, Stalnaker, and Perry to the informationalist picture: The super cool, very powerful tools and frameworks we get from their work are special cases of the more general phenomenon, of (a) leveraging information out of an observable event plus some background assumptions, and (b) exploiting what we know about the background assumptions other people are going to bring to bear in extracting information from our actions to communicate stuff to them (and exploiting what we know about

what speakers know about the background assumptions we're bringing to bear to extract information about their communicative intentions, and exploiting what we know about what hearers know about what we know about the background assumptions they're bringing to bear in order to layer on additional messages, etc.).

Consider finally (16):

16. Alpie believes that the fact that Carrie will be in the Palmer House Hilton ballroom ten minutes after the time of Alpie's utterance is relevant to some goal or project that Alpie and Betty are pursuing.

Update with (16) is supported by almost the same information that supported (15), with the difference that the assumption that Alpie is being relevant is replaced with the assumption that Alpie *believes* that he's being relevant. This is an update that's different from standard implicatures twice over, since (a) it's generated by a PIU, rather than the proposition semantically expressed, and (b) it relies not on cooperativity itself but on attempted cooperativity. Given those complications, it would be easy to expect—if one was not thinking of communication in an informationalist way—that the idea of Betty updating with (16) should come off as extremely exotic, an epicycle of an epicycle. But in thinking about actual conversational dynamics, the content of (16) in fact doesn't seem peculiar or surprising at all. Once it's pointed out, it's a completely natural, obvious, and potentially quite useful kind of inference. (Cases in which Archie has a full Gricean communicative intention toward (16) probably *are* pretty peculiar. But cases in which Betty in fact updates with (16), or in which Alpie predicts that Betty will update with (16) are perfectly ordinary, and cases in which Alpie intends Betty to update with (16), without having all of the required layers of Gricean reflexive intention, are also pretty unremarkable.)

The moral of this subsection is this: Truth is important in communication, but it's a special case of a more general phenomenon. We can extract a lot of information from peoples' speech acts by holding fixed that they're speaking truly. We can communicate a lot of stuff to other people by relying on them to hold fixed that we're speaking truly. We can establish a lot of common ground by relying on them to rely on us to rely on them etc. But truth is just one of many features we can assume an utterance has. We can extract a lot of information, communicate a lot, and establish a lot of common ground based on [full semantic information + truth], and also

from [partial semantic information + truth]. We can extract even more information, communicate more, and establish more common ground, based on [full semantic information + features other than truth], and also from [partial semantic information + features other than truth].

4.7 Sharpening and metalinguistic negotiation

Here is a justly famous example from Chris Barker (Barker, 2002):

Imagine that we are at a party. Perhaps Feynman stands before us a short distance away, drinking punch and thinking about dancing; in any case, the exact degree to which Feynman is tall is common knowledge. You ask me what counts as tall in my country. “Well,” I say, “around here, ...” and I continue by uttering [“Feynman is tall”]. This is not a descriptive use in the usual sense. I have not provided any new information about the world, or at least no new information about Feynman’s height. In fact, assuming that tall means roughly ‘having a maximal degree of height greater than a certain contextually-supplied standard’, I haven’t even provided you with any new information about the truth conditions of the word tall. All I have done is given you guidance concerning what the prevailing relevant standard for tallness happens to be in our community; in particular, that standard must be no greater than Feynman’s maximal degree of height.

This phenomenon, of using an utterance to communicate information about the resolution of context-dependence rather than to communicate the proposition expressed, is extremely friendly to goatist interpretation. The update is a kind of PIU, like many of the examples we’ve seen already. If the audience knows that Feynman is 6’1”, then they can extract the information *that ‘tall’, in the present context, picks out a degree of height no greater than 6’1”* based on [utterance + PSI + truth + Feynman’s height]. Even if they don’t know Feynman’s height, they can extract, on the basis of [utterance + PSI + truth], *that ‘tall,’ in the present context, picks out a degree of height no greater than Feynman’s*. The partial semantic information, in either case, is information that fixes everything except the resolution of the context dependence of ‘tall’.⁵¹

⁵¹ Here we’re aiming to assimilate the phenomenon Barker was drawing attention to into an informational framework. We aren’t aiming to assimilate Barker’s actual account. Barker takes the availability of sharpening uses to motivate a distinctive sort of semantic theory for the target expressions, which builds the tools for accounting for

In Barker's original case, it's natural to assume that there is an antecedently determinate (or close enough to determinate) fact of the matter regarding the standard for height. But the case doesn't have to work like that. It's not unusual to find a conversation in which part of what's at issue is where we should draw the line between the tall and the not-tall, or the rich and the not-rich, and the way the conversation proceeds is (at least in part) by making claims, regarding people whose height or financial situation is common knowledge, about who's tall and who's rich.

Let's say that Feynman, rather than standing around at a party, is standing around near the basketball court as Alphie and Betty begin to assemble their team. There is no determinate standard yet, and Alphie and Betty know it. When Alphie points to Feynman and says "Feynman is tall," part of what he aims to communicate is that the standard for tallness *should be no greater* than Feynman's height. This kind of usage is also friendly to a goatist analysis. How exactly it gets spelled out will depend a bit on how you like to treat cases where contextual resolutions are not fully settled. Let's suppose, for the sake of argument, that in the absence of an antecedently settled standard, the standard for 'tall,' as Alphie uses the word, is set by his intention that the standard be low enough to make his assertion true. (Alternative ways of setting up the case should allow for alternative, but equally goat-friendly, analyses of the case.)

If things go smoothly, Betty will update with the information *that 'tall,' in the present context, should pick out a degree of height no greater than Feynman's*. To extract this information, it is not quite enough for her to update with [utterance + PSI + truth]. In their context, where it is common knowledge that the contextual standard is not yet settled, [utterance + PSI + truth] gets Betty to the descriptive information that Alphie, in fact, uses 'tall' to pick out a degree of height no greater than Feynman's. To get from that descriptive claim about Alphie's usage to a normative claim about how the context-sensitivity of 'tall' is best resolved in their context requires that Betty take Alphie to be making *good choices* about how he uses the word 'tall'. In a Gricean spirit, let's call this assumption—the assumption that the speaker uses context-sensitive or otherwise indeterminate expressions in a way that's appropriate to the context—*metalinguistic quality*, or MQ. So Betty can extract the normative, semantic

the sharpenings into the semantics. On the sort of informational account we are suggesting, we won't add any bells or whistles to the semantics, but will instead take the sharpening uses to be cases in which among the speaker's communicative aims is bringing about a PIU with semantic information that constrains the range of available sharpenings.

information Alphie means to communicate—*that it would be good, for present purposes, to use ‘tall’ in such a way that Feynman counts as tall*—on the basis of [utterance + PSI + truth + MQ].

Of course, the question of how to use the word ‘tall’ is not, ultimately, what interests Alphie and Betty. They’re interested in assembling a successful basketball team. Here too, the goatist framework is useful in clarifying the relevant inferences. Alphie intends to communicate that ‘tall’ should be used in a certain way, not out of an interest in semantics, but because he thinks people of a certain height should be considered tall enough to be categorized together for certain purposes when it comes to assembling the team. (He and Betty might agree, for example, that only “tall” players should play Center.) In a case like this, Betty is able to extract not only the semantic information about the word ‘tall,’ but the corresponding, worldly, non-semantic information about the team. If the communication goes smoothly, she could extract, for example, the information that *players of at least Feynman’s height are tall enough to play Center*. To get to this update, Betty needs the information she used for the previous update—[utterance + PSI + truth + MQ]—plus some basic assumptions about what metalinguistic quality would amount to in this specific context. For example, that in this context a good standard for ‘tall’ will correspond to a good standard for playing Center. Call this a *linking assumption*, or LA. So, when the assumption linking the use of some expression to the aims of the conversation is added to the other information—in other words, on the basis of [utterance + PSI + truth + MQ + LA]—Betty can update with *players of at least Feynman’s height are tall enough to play Center*.

That may look like a lot of background beliefs. But recall the observation, from Section 4.1, that successful communication requires a lot of background beliefs even in the case of paradigmatic literal assertion. As we’ve emphasized, the notion of belief we have in mind allows for a lot of tacit beliefs. And crucially, a central element of the informationalist agenda is to encourage the careful elaboration of the kinds of information that would be necessary to license various updates. It should come as no real surprise if real-world communication ends up requiring the mobilization of a range of different kinds of assumptions, once you start carefully listing them out. Compare the simple case of implicature discussed in Section X.X. For Betty to extract *Alphie can’t go to the movie on Thursday night* on the basis of his utterance “I have an exam on Friday morning” would require, in the typical case, that she update on the basis of [utterance + FSI + truth + cooperativity (specifically, relevance) + information about what relevance requires in this context]. Other than the fact that the implicature relies on a FIU, rather

than a PIU, it's not really very different from the basketball update above. In other words, sharpening uses, normative sharpening uses, and even the worldly updates that go along with normative sharpening uses, are not necessarily any more exotic or demanding than other, more familiar forms of pragmatic communication.

It isn't possible here to work through all the possible variations and continuations of cases like these. But a few points are worth noting and can be made quickly.

First, we've focused here on a case involving a contextually determined threshold for a gradable adjective. But there's no reason to think that processes like these would be limited to that specific aspect of gradable adjectives. Other adjectives require not just a standard, but a specification of the scale that such a standard would relate to. Kennedy (2007) gives the example "Mexico City is larger than Tokyo," which is true if the scale for 'large' is geographic size, and false if the scale is population. And in the case of multidimensional gradable adjectives, there is the question not just of scales, and of thresholds, but of how the various dimensions should be weighted against each other to produce the scale on which a threshold can be specified. Each of these factors represents some element of what has to be known for a listener to have full semantic information and to carry out a traditional FIU. Correspondingly, any of these factors, if they are missing, would leave a listener with partial semantic information, which they could bring to bear in the form of a PIU. This in turn opens up the full range of sharpening, normative sharpening, and sharpening-related non-semantic updates discussed above, for any of these elements of context sensitivity.

Second, we see no reason to limit the phenomenon to traditional context sensitivity. Many of the PIUs discussed above, e.g. in Section X, relate not to context sensitivity, but to conventional, or ostensibly context-invariant, elements of word meaning. The "ophthalmologist is an eye doctor," foreign-language instruction, and technical term examples from Section X all relate to semantic information that is extracted by a listener, via partial information update, relating to standing or conventional word meaning. In other words, though they are instances of the same or very similar processes, these PIUs relate to information not about semantically-relevant aspects of the context, but rather to the *character*, in Kaplan's sense, of the relevant expressions. This means that for the goatist, there is no deep difference between a sharpening usage of 'tall,' or 'rich'—meant to convey information or proposals related to traditional context sensitivity—and a sharpening usage of 'torture,' 'sandwich,' or 'planet'—

meant to convey information or proposals related to the disambiguation, precisification, sharpening, or modulation of word meaning in some broader sense.⁵²

Finally, and quite generally, if speakers are able to communicate a certain kind of information, then speakers and their audiences should be able to express disagreements about that information. And so, when it comes to sharpening uses and their associated non-semantic updates, we take goatism to provide a natural framework for understanding the phenomenon of *metalinguistic negotiation*.⁵³ Plunkett and Sundell define *metalinguistic disputes* as disputes over the correctness or appropriateness of opposing sharpening claims. And when those claims are naturally seen as opposing bids in a process meant not to describe but to shape the semantic facts, they describe the disputes as *metalinguistic negotiations*.

While we have not in general focused on updates in the context of disagreement or debate, for the goatist this kind of phenomenon is not a surprising or exotic development, but a completely predictable outgrowth of the same, ubiquitous communicative mechanisms at work in standard implicatures, PIUs of the kind Stalnaker and Perry draw attention to, and even straightforward assertion. Speakers often rely on audiences to extract, via partial information update, claims about word meanings or semantically relevant aspects of the context. Audiences sometimes accept and update with those claims. But often enough, they instead (a) back off from assumptions of *truth* or *metalinguistic quality* to assumptions of *sincerity* or *attempted-quality*, (b) update with the fact that the speaker believes the relevant claim—but not with the claim itself, and (c) go on to express, by equivalent mechanisms, their own, alternative views. When they do so, the parties to the conversation are engaged in a metalinguistic dispute, and when the claims are normative, the metalinguistic dispute takes the form of a metalinguistic negotiation.

Conclusion

All update is goat update. (Or a downstream consequence of goat update.)

⁵² This approach is in contrast to analyses that take the phenomenon to be real, but to be limited to context-sensitive expressions or even a special class of context-sensitive expressions. As noted above, (Barker, 2002, 2009, 2012) offers an account of sharpening uses of gradable adjectives, but does not suggest that it be extended to expressions that are not semantically context sensitive. (However, we take it that there's nothing in his account that would prevent it from being extended in this way.) Justin Khoo (Khoo, 2020) offers an analysis of this kind of phenomenon as it occurs with certain indexicals like "rich," which he takes to constitute a distinctive subclass of context dependent expression—what he calls *quasi indexicals*.

⁵³ (Plunkett & Sundell, 2013, 2019)

References

- Almotahari, M., & Glick, E. (2011). Context, Content, and Epistemic Transparency. *Mind*, 119(476), 1067–1086. <https://doi.org/10.1093/mind/fzq087>
- Armstrong, J. (2018). Provincialism in Pragmatics. *Philosophical Perspectives*, 32(1), 5–40. <https://doi.org/10.1111/phpe.12114>
- Armstrong, J. (2023). Communication before communicative intentions. *Noûs*, 57(1), 26–50. <https://doi.org/10.1111/nous.12396>
- Barker, C. (2002). The dynamics of vagueness. *Linguistics and Philosophy*, 25(1), 1–36.
- Barker, C. (2009). Clarity and the grammar of skepticism. *Mind & Language*, 24(3), 253–273.
- Barker, C. (2012). *Negotiating taste*. 1–19.
- Barwise, J., & Perry, J. (1983). *Situations and Attitudes* (Vol. 78, Issue 11, pp. 668–691). MIT Press.
- Bowker, M. (2019). Saying a bundle: Meaning, intention, and underdetermination. *Synthese*, 196(10), 4229–4252. <https://doi.org/10.1007/s11229-017-1652-0>
- Brandom, R. (1983). Asserting. *Noûs*, 17(4), 637–650.
- Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment* (Vol. 183, Issue 183, p. 238). Harvard University Press.
- Brown, P., & Levinson, S. (1978). *Politeness: Some Universals in Language Use*. Cambridge University Press.
- Buchanan, R., & Ostertag, G. (2005). Has the Problem of Incompleteness Rested on a Mistake? *Mind*, 114(456), 889–913. <https://doi.org/10.1093/mind/fzi889>
- Camp, E. (2013). Slurring Perspectives. *Analytic Philosophy*, 54(3), 330–349.

<https://doi.org/10.1111/phib.12022>

Camp, E. (2017). Why metaphors make good insults: Perspectives, presupposition, and pragmatics. *Philosophical Studies*, 174(1), 47–64.

<https://doi.org/10.1007/s11098-015-0525-y>

Camp, E. (2018). Insinuation, Common Ground, and the Conversational Record. In D. Harris, D. Fogal, & M. Moss (Eds.), *New Work on Speech Acts* (pp. 40–66). Oxford University Press.

Camp, E., & Nowak, E. (n.d.). *Linguistic Variation, Agency, and Style* (U. Stojnic & E. Lepore, Eds.).

Caponetto, L. (2020). Undoing things with words. *Synthese*, 197(6), 2399–2414.

<https://doi.org/10.1007/s11229-018-1805-9>

Caponetto, L., & Cepollaro, B. (2022). Bending as Counterspeech. *Ethical Theory and Moral Practice*. <https://doi.org/10.1007/s10677-022-10334-4>

Chomsky, N. (1986). *Knowledge of Language: Its Nature, Origins, and Use*. Praeger.

Davidson, D. (1968). On saying that. *Synthese*, 19(1–2), 130–146.

<https://doi.org/10.1007/BF00568054>

Davidson, D. (1986). A Nice Derangement of Epitaphs. In E. LePore (Ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson* (pp. 433–446). Blackwell.

De Ponte, M., Korta, K., & Perry, J. (2023). Critical Pragmatics: Nine Misconceptions. *Topoi*, 42(4), 913–923. <https://doi.org/10.1007/s11245-023-09903-8>

Dretske, F. (1981). *Knowledge and the Flow of Information* (Vol. 92, Issue 3, pp. 452–454). MIT Press.

Dretske, F. (1995). *Naturalizing the Mind* (Issue 4, pp. 528–537). MIT Press.

Egan, A. (2009). Billboards, bombs and shotgun weddings. *Synthese*, 166(2), 251–279.

- Field, H. H. (1978). Mental representation. *Erkenntnis*, 13(1), 9–61.
<https://doi.org/10.1007/BF00160888>
- Fricker, E. (1994). Against Gullibility. In A. Chakrabarti & B. K. Matilal (Eds.), *Knowing from Words*. Kluwer Academic Publishers.
- Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford University Press.
- Friedman, J. (2013). Question-Directed Attitudes. *Philosophical Perspectives*, 27(1), 145–174.
<https://doi.org/10.1111/phpe.12026>
- Gendler, T. S. (2011). On the epistemic costs of implicit bias. *Philosophical Studies*, 156(1), 33–63. <https://doi.org/10.1007/s11098-011-9801-7>
- Gilbert, D. (1991). How mental systems believe. *American Psychologist*, 46(2), 107–119.
- Gilbert, D., Krull, D., & Malone, P. (1990). Unbelieving the unbelievable: Some problems in the rejection of false information. *Journal of Personality and Social Psychology*, 59(4), 601–613.
- Gilbert, D., Tafarodi, R., & Malone, P. (1993). You can't not believe everything you read. *Journal of Personality and Social Psychology*, 65, 221–221.
- Grice, H. P. (1975). Logic and Conversation. In Cole, P (Ed.), *Syntax and Semantics 3: Pragmatics*. Academic Press.
- Groeneveld, W., & Veltman, F. (1997). An update on “might.” *Journal of Logic*.
<http://www.springerlink.com/index/Q78602641N341U7W.pdf>
- Hardwig, J. (1985). Epistemic Dependence. *Journal of Philosophy*, 82(7), 335–349.
- Harris, D. W. (2020). We talk to people, not contexts. *Philosophical Studies*, 177(9), 2713–2733.
- Hawthorne, J., & Magidor, O. (2009). Assertion, Context, and Epistemic Accessibility. *Mind*, 118(470), 377–397. <https://doi.org/10.1093/mind/fzp060>
- Israel, D., & Perry, J. (1990). What is Information. In P. Hanson (Ed.), *Information, Language*

- and Cognition* (pp. 1–19). University of British Columbia Press.
- Jackson, F. (2002). Critical Notice of *Knowledge and Its Limits* by Timothy Williamson. *Australasian Journal of Philosophy*, 80(4), 516–521.
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–566). Oxford University Press.
- http://scholar.google.com/scholar?q=related:Lfif5IFS8SUJ:scholar.google.com/&hl=en&num=30&as_sdt=0,5
- Kaplan, D. (ms). *The Meaning of “Ouch” and “Oops.”* 1–29.
- Khoo, J. (2020). Quasi Indexicals. *Philosophy and Phenomenological Research*, 100(1), 26–53.
- <https://doi.org/10.1111/phpr.12519>
- King, J. C. (2021). *Felicitous Underspecification: Contextually Sensitive Expressions Lacking Unique Semantic Values in Context*. Oxford University Press.
- Kirk-Giannini, C. D. (2019). Why horizontalism. *Philosophical Studies*.
- <https://doi.org/10.1007/s11098-019-01350-9>
- Kirk-Giannini, C. D. (2024). Do Not Diagonalize. In E. Lepore & U. Stojnic (Eds.), *The Oxford Handbook of Contemporary Philosophy of Language*. Oxford University Press.
- Korta, K., & Perry, J. (2011). *Critical pragmatics: An inquiry into reference and communication*. Cambridge University Press.
- Kukla, R. (2014). Performative Force, Convention, and Discursive Injustice. *Hypatia*, 29(2), 440–457. <https://doi.org/10.1111/j.1527-2001.2012.01316.x>
- Kukla, R., & Lance, M. N. (2009). *“Yo!” and “Lo!”: The pragmatic topography of the space of reasons*. Harvard University Press.
- Lackey, J. (1999). Testimonial Knowledge and Transmission. *The Philosophical Quarterly*, 49(197), 471–490. <https://doi.org/10.1111/1467-9213.00154>
- Lackey, J. (2006). Learning from Words. *Philosophy and Phenomenological Research*, 73(1),

- 77–101. <https://doi.org/10.1111/j.1933-1592.2006.tb00605.x>
- Lackey, J. (2008). *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press.
- Langton, R. (2017). Blocking as counter-speech. *New Work on Speech Acts*.
http://www.academia.edu/download/50922808/Blocking_as_counter-speech_Dec_2016.pdf
- Lewis, D. (1966). An argument for the identity theory. *The Journal of Philosophy*.
<http://www.jstor.org/stable/2024524>
- Lewis, D. (1980). Index, context, and content. In S. Kanger & S. Öhman (Eds.), *Philosophy and Grammar* (pp. 79–100). Reidel.
- Lewis, D. (1983). *Mad Pain and Martian Pain*. Oxford University Press.
<https://doi.org/10.1093/0195032047.003.0009>
- Mandelbaum, E. (2011). *The architecture of belief: An essay on the unbearable automaticity of believing* [PhD Thesis, University of North Carolina].
<http://gradworks.umi.com/34/08/3408816.html>
- Mandelbaum, E. (2014). Thinking is believing. *Inquiry*, 57(1), 55–96.
<https://doi.org/10.1080/0020174X.2014.858417>
- Murray, S. E. (2014). Varieties of update. *Semantics and Pragmatics*, 7.
<https://doi.org/10.3765/sp.7.2>
- Perry, J. (1988). Cognitive significance and new theories of reference. *Noûs*, 22(1), 1–18.
- Perry, J. (2001). *Reference and Reflexivity* (Issue 445). Center for the Study of Language and Inf.
- Plunkett, D., & Sundell, T. (2013). *Disagreement and the Semantics of Normative and Evaluative Terms*. 13(23), 37.
- Plunkett, D., & Sundell, T. (2019). Metalinguistic negotiation and speaker error. *Inquiry*, 1–26.

<https://doi.org/10.1080/0020174X.2019.1610055>

Predelli, S. (2013). *Meaning Without Truth*. Oxford University Press.

Ramírez, E. G. (2010). Wookiee Statements, Semanticism, and Reasonable Assertion. *Revista de Filosofía*, 35.

Rescorla, M. (2009). Assertion and its constitutive norms. *Philosophy and Phenomenological Research*, 79(1), 98–130.

Roberts, C. (2004). Context in dynamic interpretation. *Handbook of Contemporary Pragmatic Theory*, 197–220.

Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5, 1–69. <https://doi.org/10.3765/sp.5.6>

Rudolph, R. E. (2021). Conceptual exploration. *Inquiry*, 1–26.

<https://doi.org/10.1080/0020174X.2021.2002053>

Rudolph, R. E., & Kocurek, A. W. (2021). Comparing Conventions. *Semantics and Linguistic Theory*, 30, 294. <https://doi.org/10.3765/salt.v30i0.4820>

Rudolph, R. E., Kocurek, A. W., & Jerzak, E. (2020). Against Conventional Wisdom. *Philosophers' Imprint*, 20(22), 1–27.

Schwitzgebel, E. (2002). A phenomenal, dispositional account of belief. *Noûs*, 36(2), 249–275.

Stalnaker, R. (1978). Assertion. *Syntax and Semantics*, 9(Pragmatics), 315–332.

Stalnaker, R. (1984). *Inquiry*. MIT Press.

Stalnaker, R. (1987). Semantics for belief. *Philosophical Topics*.

http://www.pdcnet.org/philtopics/content/philtopics_1987_0015_0001_0177_0190

Stalnaker, R. (1998). On the representation of context. *Journal of Logic*.

<http://www.springerlink.com/index/q6r5g8808715p200.pdf>

Stalnaker, R. (2002). Common ground. *Linguistics and Philosophy*.

<http://www.springerlink.com/index/v20u533483228545.pdf>

- Stalnaker, R. (2006). Assertion revisited: On the interpretation of two-dimensional modal semantics. *Two-Dimensional Semantics*.
<http://books.google.com/books?hl=en&lr=&id=I2ATDAAAQBAJ&oi=fnd&pg=PA293&dq=assertion+revisited&ots=V39DQ6gEjH&sig=KK2BzAqFD54Pkia2zq5VbN7XeLQ>
- Stalnaker, R. (2009). On Hawthorne and Magidor on Assertion, Context, and Epistemic Accessibility. *Mind*, 118(470), 399–409. <https://doi.org/10.1093/mind/fzp061>
- Stalnaker, R. (2014). *Context*. Oxford University Press, USA.
http://books.google.com/books?id=zlrBAwAAQBAJ&printsec=frontcover&dq=intitle:Context+inauthor:stalnaker&hl=&cd=1&source=gbs_api
- Sundell, T. (2009). *Conflict and Content (PhD Dissertation)*. University of Michigan.
- Sundell, T. (2014). Aesthetic Negotiation. In *The Semantics of Aesthetic Judgment* (pp. 1–19).
- Sundell, T. (2016). The tasty, the bold, and the beautiful. *Inquiry*, 59(6), 793–818.
<https://doi.org/10.1080/0020174X.2016.1208918>
- Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*.
<http://www.springerlink.com/index/FKP72335W64G0131.pdf>
- von Fintel, K., & Gillies, A. (2009). “*Might*” made right. <http://philpapers.org/rec/VONMMR>
- Waller, M. (2013). Dynamics of epistemic modality. *Philosophical Review*.
<http://philreview.dukejournals.org/content/122/1/45.short>
- Williamson, T. (2000). *Knowledge and its Limits* (Issue 297, pp. 460–464). Oxford University Press.
- Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464), 983–1026.