

Deliberating, Concluding, and Entailing

Jim Pryor

5 Jan 2012

We will be considering the three topics of my title—or as we might also call them, reasoning, rationally responding, and logic. In many philosophers’ minds these are loosely but firmly connected. Too firmly. It’s not easy to identify a rigorous thesis they definitely accept and I definitely reject. But I will be urging these three notions are farther apart and explanatorily more independent than is usually assumed. Most notably, I’ll be opposing “Closure Principles” for reasonable belief (in ways that I think don’t depend on the details of how such principles are formulated).

This paper is too telegraphic, too heavily qualified and unnecessarily complex, and too apologetic to boot! Sorry. My first drafts tend to come out this way, despite my better intentions. I take comfort at least in the thought that there are interesting issues here (no doubt too many of them), and that this might conceivably make for good discussion.

Some of what I’ll be saying may remind you of Harman’s *Change in View*. That association is in places appropriate. But I’ll make no attempt to separate where I see myself as defending Harman, where extending, and where diverging. It will be more effective to make the case from scratch in my own terms.

1. Deliberating

I’ll begin with the notion of deliberating or reasoning. Or rather, I’ll begin with what these terms mean in English. What philosophers almost always associate them with is something else, perhaps overlapping but importantly different. That’s what I’ll instead call “concluding” or “rationally responding.” It will be the topic of the next section.

So what do I mean by “deliberating”?

You’re trying to solve a logic puzzle, or trying to evaluate and rank a grad student application. These may involve some important non-discursive elements, but at the same time they are paradigms of the activities folk call “reasoning.” Consider the narrative of your stream of consciousness while you engage in

these activities. That would look nothing like a logical proof—not even a proof with suppressed premises. It proceeds in fits and starts. There’s backtracking, sudden shifts in strategy, and idle mulling about. Some lines of thought are abandoned unfinished. You sometimes send off background mental searches—“What counterexample does this remind me of?”—and think about other things until the result comes back. All of this makes up the occurrent directed thinking of someone who’s epistemically healthy, and may be reasoning excellently. It is good reasoning.

Skill in this activity requires much more than just knowing how to follow a chain of entailments to their bitter end. In fact, that’s probably only a small part of the skill. Consider that teaching a child or an AI to do this well will largely be a matter of identifying techniques and heuristics that look nothing like a logical formalism. Consider also how some mathematicians may be better than their colleagues at quickly figuring out the interesting properties of a new construction. Yet their knowledge of background mathematical facts may be no better than the others. And each may be equally capable of following a given entailment. The better ones may just be more skilled at figuring out where to direct their efforts. This makes them better mathematical reasoners.

Moreover, these activities we call deliberating or reasoning vaguely blur into other inquisitive activities, like hunting down additional evidence, choosing which of your old notes to go back and review, and so on. (The skills Sherlock Holmes was better in than his brother Mycroft.)

In all these ways, deliberating or reasoning is quite unlike the model we encourage when introducing students to formal logic: a matter of proceeding step by logical step from premises to conclusion. That model may be pedagogically useful for the student, but it’s a serious handicap for the theorist.

2. “Concluding”

Closer to that model is what I’d like to call “concluding” or “rationally responding” (and other philosophers will simply call “reasoning”). Though the label “concluding” is misleading in several ways. And even once those distractions are bypassed, this notion will turn out to importantly diverge from the model of the logic classroom, too.

But first, what is it we’re now talking about? And what’s its relation to the activities discussed in the previous section?

By “concluding” I here mean the reason-guided transition from current attitudes to further attitudes or actions—“guided” in the sense of being *a response to* those reasons or considerations. As the illustrations we considered above make clear, this is only a small part of what goes on in the thinking processes that are folk paradigms of reasoning well. It is more like an end-goal of those processes. That makes the term “concluding” sound well-chosen. But on the other hand,

these transitions needn't be conclusive. The chain of reasoning may continue on wrestling with other questions. And the transition need not be, and needn't be thought by us to be, a response to conclusive reasons. Finally, and most interestingly, the notion of "concluding" suggests a kind of "settling," which is an appropriate image only for some of the doxastic and epistemic notions we might theorize about. We'll come back to this in the next section.

In fact, I'm calling this notion "concluding" mainly so that we can attend to all the ways in which that label misleads.

It's easier to see the difference between what I'm calling "deliberating" and what I'm calling "concluding" in the epistemic realm. The mental activities of deliberating are presumably intentional and themselves responsive to reasons. If you know that thinking about such-and-such an objection makes you impatient, and you judge you shouldn't get impatient with this audience, you would reasonably attempt to direct your attention to other aspects of the dialectic you're engaged in. The reasons involved here will be *practical ones* (even if they might sometimes be connected to goals special to inquiry), and the responsiveness you do or should have to those reasons would be part of practical rationality. Whereas the reasonable transitions you're working your way towards would be epistemic responses to epistemic reasons.

I suspect that the rationality of deliberating about Y-ing and the rationality of Y-ing can *also* be usefully separated even when Y-ing is itself a practical activity. But since in that case the deliberating and the Y-ing are both practical, it's more contentious and harder to think about.

Now you can deliberate or "reason from" things you don't believe. Similarly, you can "reason with" rules you don't accept. All this means you can engage in the activity discussed in section 1, even to what feels like a satisfying conclusion, without meaning to reach thereby any new belief of your own. You needn't even reach any *conditional* belief—that if these starting assumptions are accepted this conclusion follows. For you may know you've lost track of some of the optional but reasonable auxiliary assumptions you relied on along the way, and yet have reasoned competently for all that.

Similarly, you can reason to things you reasonably find unacceptable from things you (still) reasonably accept. It's an interesting question what your rational options then look like, but it's doubtful that you always must believe the conclusion.

Finally, it's an open possibility, I think even a plausible one, that there can be procedures for deliberately "making up your mind" or forming belief that aren't rightly counted as "reasoning" or "deliberation." For instance, someone might have the doxastic policy of "Think what Nietzsche would!" and have mental rituals for determining what that is that aren't a matter of reasoning it out. They might be more like what an artist does when staring at a blank canvas. Another might have the policy of always going with their second thought.

So not only does reasoning in the folk sense not have to end in a rational transition: neither does the latter need to be prefaced by the former. Of course, many would already accept this last point since they think some beliefs can be reasonable yet spontaneous. We're now making the additional point that even when beliefs are formed deliberately and non-spontaneously, the process by which one does so need not necessarily count as "reasoning" or "deliberating" in the sense of section 1.

(It's one thing to say that something can be a rational transition, and a stronger claim to say that it's a *good or reasonable* rational transition. Here I'm only making the more modest claim that deliberate procedures can end, and be meant to end, in rational transitions, without being "reasoning" or "deliberating." Not the more ambitious claim that such procedures can possibly end in *good* rational transitions.)

I think it's a shame that philosophers have co-opted the folk term "reasoning" in the way they have, for those segments of our thinking that are (or purport to be) reasonable transitions to new doxastic attitudes, or new choices of what to do. The broader activity we discussed in section 1 is also very interesting to our fields, and deserves a name, and in plain English it already had one.

However, for most of what follows, we will *not* be concerned with those broader activities (except at one point where I'll remind you of them). Instead, we'll be focusing on the mental processes introduced in this section. Our concern will be: what are the relations between "concludings," or reasonable transitions to new doxastic attitudes or choices, on the one hand, and on the other hand the logical relations among the contents of the attitudes we transition between?

A popular view is that some logical relations suffice for the reasonableness of the corresponding transitions. Except this is tricky to state properly: we have to acknowledge that we're not logically omniscient, that people may believe the right things for the wrong reasons, and so on. Still, many philosophers are convinced that something of this spirit—some "Closure Principle" for reasonable belief—must be correct, despite the trickiness of identifying the right letter.

I will not fuss with the letter of such proposals, either. I will try instead to foster suspicion about their spirit.

3. Chunky versus graded attitudes

Some of our doxastic and volitional attitudes are graded: we're more confident in this than that. We're less prepared to believe one way or another about something else. (Perhaps there are practical analogues, but I won't argue for that.)

On the other hand, other attitudes are "chunkier" or more discrete. Many (perhaps all) of these attitudes invite the language of "settling" on a view or a

choice, or “making up your mind.” Doxastically, you might settle on a verdict of “yes” or “no”, or you might instead settle for the time being on agnosticism. Practically, you settle on a course of action.

There are paradigm phenomenal and dispositional differences I associate with the language of “settling” or “making up your mind.” Think of the experience of *taking yourself* to have decided which of your suitors to marry, but then seeing from your continuing deliberations and actions that you haven’t in fact yet made up your mind. You’re discovering that the phenomena characteristic of having settled on a choice are in fact mostly missing.

Those phenomena have analogues in the doxastic realm. I like the image of the police filing away a crime report, perhaps with their verdict, or perhaps with a stamp of “unknown.” There are distinctive psychological markers of your having reached a settled view about a question. Your having settled and filed away your report on a question doesn’t mean you can’t be ready to respond to further evidence when you encounter it. It doesn’t even mean you must have given up on looking for evidence. You can look for evidence about questions you’ve got a settled view about (in principle, you may even do so impartially); and you may be in situations where it’s useful for you to do that. But just as there’s a recognizably different theme or feel to your thinking and actions when you’ve decided what to do, so too is there often a recognizably different theme or feel when you’ve decided yes, P must be true, or when you’ve decided to settle on withholding judgment for now about P. This psychological shift need not be the result of any new evidence or new insights. Sometimes we just make up our minds. Or find that we’ve made them made up.

What does deciding or choosing in these ways amount to? Cases we’d say exhibit it have vague overlapping similarities. Perhaps there are natural explanatory states underlying those similarities, and there’s a definite fact of the matter when they are present, and to what extent. Or perhaps it’s more like Wittgenstein encouraged us to think of “games.” What comes first are the vague overlapping similarities; we impose the homogeneous predicate of “having decided” later. In different settings, we negotiate the extension of that predicate differently, and aren’t constrained in doing so by an underlying reality that explains the patterns we see in the symptoms.

I have some sympathy for each of these pictures. But for this discussion it shouldn’t matter which way the metaphysics falls here.

On the other hand, as I said, some of our doxastic reactions are more graded.

I can take some evidence, newly acquired or newly appreciated, to support *more* confidence in P, without yet having made up my mind that P is true. I think this kind of graded doxastic responding can go on even *after* I’ve “settled” or made up my mind. Can’t one have the experience of now being even *more* sure of something one already had decided was true? But from experience, it does seem that there tends to be much less graded doxastic movement after you’ve “settled.”

Changes in graded doxastic attitudes seem to be guided by reasons, and seem to be ex post assessable, as much as changes in settled attitudes. You can become more confident for the wrong reasons, even when the right reasons are also available.

I think it's possible to fail to have any graded attitude towards a question: for example, questions you're not yet capable of understanding. Perhaps also to some questions you've never considered. (Though not having considered a question doesn't entail that you have no graded attitude towards it.) An interesting hypothesis would be that whenever you *actively consider* a question, you cannot fail to have *some* kind of graded doxastic stance towards it. Perhaps a doxastic stance that's fundamentally imprecise, but a stance nonetheless. I'm not sure whether this hypothesis is correct. At any rate, even if it turns out one can *sometimes* fail to have any such stance, there will also be many cases where one *does* have a graded doxastic stance but has yet failed to reach a "settled" view. Your confidence can wax and wane many times before you make up your mind.

At the same time, all that waxing and waning shouldn't be *identified with* all the moves in your deliberating. Many moves in your deliberating won't manifest themselves in a change of graded doxastic attitude. In the first place, as we said before, you can deliberate from starting points you don't accept. Moreover, even when you do accept the starting points, much of deliberating will be *feeling your way towards* how your doxastic responses should evolve—both graded and settled species of them. A lot of your thinking might still be feeling around, without yet having made any doxastic progress, even of the graded sort.

In addition to chunky versus graded doxastic states, we also have chunky versus graded epistemic assessments. How these interact is not entirely clear. But when we theorize about knowing or justifiably believing, we're working with chunky attitudes and chunky assessments of them. When we theorize about *how* justified you are in believing, we're working with a graded assessment of a chunky attitude. When we theorize about *having some justification to be more confident*, we're epistemically assessing a graded attitude; it's unclear whether we're doing so in a chunky or graded way.

How do chunky and graded doxastic attitudes relate to each other? These are large questions.

Often one sees it assumed that, at least in a given case, there will be a threshold of confidence such that any more confidence than that counts as believing. That's not the picture I'm working with. I think of believing as having settled, and being more confident than so-and-so is entirely compatible with not having yet settled. There are other orthogonal issues that might be bound up with believing as well: what is your attitude towards the probative value of seeking corroboration, gathering more evidence, and so on? These things aren't fixed by how confident you are about the proposition at hand. Yet we might think them relevant to whether you count as believing that proposition. Similarly, whether

you believe that P may constitutively constrain, or depend on, what ways you're prepared to deliberate, in the sense discussed in section 1. Believing P might require a readiness to "use P as a premise in your deliberations." It's far from clear that any such constitutive connections are imposed by any (non-maximal) degree of confidence.

Also, I'm not sure that every chunky doxastic attitude comes along with a graded confidence. For instance, one chunky attitude is settled agnosticism, and it's not clear what kind of graded confidence, if any, that would be constituted by. So some chunky attitudes may be possible in the absence of any graded attitudes. Or perhaps we should instead have a larger palette of what the envisaged graded attitudes may be.

As I said, these are large questions.

The many options here will complicate our discussion of our main target. I think it will be easiest to see our way forward if we focus on graded attitudes. At certain places I will point out how attending only to the chunky attitudes might mislead us.

4. Which logic is correct?

One central idea in what follows will be that you can be reasonably uncertain whether you've concluded something validly. One way that can happen is by having reasonable doubts about your performance. Another way is by having reasonable doubts about what validity amounts to. The force of these second doubts are underestimated by many philosophers, so let's take a moment to emphasize them.

Of course, your usual philosopher will acknowledge, there are parts of logic that are rightly much-disputed. Much about modals, for example. And conditionals: perhaps there are things to be reasonably disputed about what "if" in our native languages and thinking amounts to. Several philosophers have argued that our native "if" doesn't generally validate modus ponens. (Van McGee famously argued so, but a slew of recent papers have also pushed for this, for somewhat different reasons. MacFarlane and Kolodny are one example; also Yalcin; Gillies?)

And then there are the semantic paradoxes, which prompt some philosophers to go for relevance logic.

But look, if we just focus on paradox-free parts of our language, and set aside the meaning of "if", then surely there's not much room for dispute? Well, there is Dummett and his intuitionism; though that seems to rely on unpopular metaphysical commitments. But we can be generous. Intuitionistic validities are a subset of the classical ones. So with these restrictions in place, can't we at least say there's no room for reasonable doubt about the intuitionistic validities?

But yes, there is. Are all the terms in your language logically guaranteed to refer to existing objects? For example, is there a term in your language for division? And is it right for it to be a logical truth that:

1. Exists x : $x = 0/0$.

No? Well, then you'll need a free logic. And there is controversy about which free logic is correct. (If you prefer to treat "0/0" as a Russellian description, rather than a singular term, then we need a non-referring name to induce this pressure. There we confront some additional controversies.)

Well, how about propositional logics? Surely at least with the validities of intuitionistic propositional logic there's no room for rational dispute?

But yes, there is. (In fact, there's room for dispute even about its "implicational fragment," which is strictly weaker.) Is your language fully determinate and bivalent? For example, is the following a logical truth:

2. I haven't stopped swimming or I have stopped swimming.

Presumably not: say, if I never swam. And fortunately, intuitionistic logic doesn't tell us *it* is a logical truth. But it will tell us *other* things are logical truths which should provoke the same resistance that this example provokes. Unfortunately, those things will all either involve material implication or will be difficult to parse in natural language. So we can't wholly avoid thinking about conditionals—though we can avoid thinking about the meaning of our natural language "if." Do you want to count this as a logical truth:

3. (I have stopped swimming \supset I haven't stopped swimming) \supset I haven't stopped swimming.

It's tricky to evaluate this, if we don't assume " \supset " is a term of natural language. But I find it difficult to hear the claim without triggering the presupposition that I once swam. So if I never did swim, presumably the claim fails to be true. Yet it's a tautology even of a logic strictly weaker than intuitionistic.

The natural thing to say about such sentences is that our language isn't bivalent: that, when I never swam, "I stopped swimming" isn't true but neither is it false. And though intuitionistic logic rejects excluded middle, it's not well understood as rejecting bivalence. At any rate, it's not a natural vehicle for rejecting bivalence. If we want to allow some sentences to be indeterminate or have undefined truth values, then we again face a variety of logical choices and it's again controversial which logic is correct.

Or should we rather say that "I haven't stopped swimming" isn't strictly the negation of "I have stopped swimming"? Then it becomes less obvious what

the logical structure of our native talk and thinking really is, and how many of the entailments of the logics we're considering they exhibit.

If one takes these lessons seriously, there should be plenty of room for competent reasonable doubt about which logic is correct. It's unclear what the right logics are, even restricted to the intended domains of propositional and predicate logic. (Their *intended* domains are fragments of reasoning we actually engage in. They may be correct mathematical theories of their *stipulated* domains, which are artificial languages having certain formal properties.)

As I'll argue in the next section, there is even more room for *incompetent* reasonable doubt. That is, one can reasonably be moved by arguments that are in fact flawed. And there might be many such compelling arguments for logics that aren't in fact correct, or against inferences that are in fact validated by the correct logic. But even restricting ourselves to unflawed arguments, we still should allow that there are now reasonable grounds for doubt about which logic is correct. And it's not the case that all the viable competitors validate some common core. So we don't yet seem to be in a position to say *with certainty* about any logical rule that it is in fact generally correct.

(What does "correctness" amount to? Another large question. Many understand this in terms of truth-preservation, though the story becomes more complicated when we're thinking non-bivalently. Hartry Field recently argued that we instead must begin with a more primitive notion of "following." I have sympathy for his claims. But I mention them here only to register that I am not yet assuming any specific rigorous unpacking of "correct.")

Well, perhaps you'll allow that much: we can't reasonably be *certain* about interesting disputes in logic. But still, we assume, *some* logic will be correct whether we know it or not. And you might think, if your beliefs transition in ways that are *in fact* valid—and not by accident, let's suppose we throw in as well that you correctly do take those ways to be valid—then that must suffice for your transition to be a reasonable one? (At least, if the beliefs you started with were reasonable. I assume this throughout.)

But no, I don't think it must suffice. For if reasonable doubts about logic are possible, shouldn't they be possible about the transition you're then making? Even if it is a valid one, and is correctly taken by you to be so, can't you be in a position where you should (and perhaps even do) have some doubts about it? If so, then even when such grounds for doubt are absent, can it be the validity alone that makes you justified? Won't the absence of reasons to doubt the validity also be playing a role?

We'll look at such cases more closely in the next section. I think of them as posing the question "What are the epistemic effects of bad objections?" The person offering the objection might be doing so reasonably (or might only be a hypothetical person). But the objection is "bad" in that it's saying, incorrectly, that a transition which is in fact logically correct isn't.

There's a related question it's also useful to consider, that I think of as "What are the epistemic effects of bad arguments?" That is, arguments that aren't in fact valid but strike us as being so. We'll take this question up first, then turn to the other.

5. Bad arguments and bad objections

Consider a pair of logicians. They know themselves to have equally good (and pretty good though not perfect) track-records about logical questions of the sort they're now considering. The first logician, Zachary, reasons impeccably to the conclusion that P. The second logician, Yasmine, reasons herself to the same conclusion, but in fact her reasoning is flawed in a way that's not now apparent to her. Imagine that Yasmine's conclusion strikes her with the same sort of subjective compellingness that it strikes Zachary—at least, to the extent that this is compatible with the one having reasoned impeccably and the other having made a subtle mistake.

How does Yasmine's epistemic standing towards P compare to Zachary's?

One answer, which doesn't satisfy me, is that their justification to believe P is equally good. I'll call that the "subjectivist" answer.

On the other hand, I'm also not satisfied with the idea that Yasmine has no justification at all to believe P, but only seems to. Surely it's more reasonable for Yasmine to believe P than it would be for someone who had not even the appearance of a proof, and was just guessing.

This last idea needs to be handled carefully. It'll help to bring in a third reasoner, Xavier. Xavier arrives at the belief that P via his crystal ball. As it happens, his crystal ball has heretofore had as good a track record as Zachary's and Yasmine's actual track record at doing logic. That doesn't have to be too fabulous. As I said, they have made occasional mistakes, like the rest of us. They can be competent logicians for all that. And we don't have to suppose the crystal ball really is magical, just that it's in fact happened to have a good track record to this point. There are plenty of details I've left unspecified. But I can imagine them filled in such a way that, even if crystal balls *aren't* good authorities about logic, it can still be *somewhat* reasonable for Xavier to believe P in response to his crystal's say-so. At least, *somewhat more* reasonable for him to believe that than if the crystal hadn't said so. He does after all have a decent track-record to this point.

And now returning to Yasmine, the anti-subjectivist should of course allow her at least the same standing. It's not that she's got *no justification at all* to believe P. How often is she wrong when she does some logic and the conclusion she reaches seems that compelling? Granted, sometimes, but not that often. So the worst we should say about Yasmine is that she does have that much, and that kind of, justification, but no more.

Yet this too is still unsatisfying. It doesn't seem right to me to say that Yasmine is no better off epistemically towards P than Xavier is. Even if it's agreed that that's somewhat better than off than just guessing.

What should we say?

What I'd like to say is that Yasmine is worse off than Zachary in some ways, but not so worse off as Xavier, who is himself better off than a mere guesser. Xavier has some baseline inductive grounds for believing P that the others share. But intuitively Yasmine and Zachary have more. Yasmine's flawed reasoning—her “bad argument”—makes P more reasonable for her. But equally intuitively, she's not on fully equal terms with Zachary, either. One crude way this might be developed is to say that Xavier has reason to be confident to degree x in P, Yasmine to degree y , and Zachary to degree z , where $0.5 < x < y < z$. As we'll see, I don't in the end think this is the best way to understand what's going on. But it gives us a proposal to work with.

That's an initial stab at the question “What are the epistemic effects of bad arguments?”—that is, arguments that appear valid but aren't. Next we'll consider the question, “What are the epistemic effects of bad objections?”—that is, when some evidence challenges a piece of reasoning that was in fact valid.

Recall that Zachary reasoned impeccably. Say that happened in the morning. In the afternoon, his colleagues come over and review his work. They present him with good evidence to believe he's made a mistake. He hasn't made a mistake; but there might of course be good evidence that he did. Perhaps he used a proof technique that his colleagues know scholars with his limited experience tend to mess up. Or perhaps he's in a small airplane at sufficient altitude to induce hypoxia, which affects one's reasoning abilities in ways one's not then able to detect. (David Christensen has examples like that.)

Or, a different variation: in the evening, Zachary's colleagues come over and tell him he's relied on a logical rule that isn't generally valid, and that the issues he's working with are some of the vulnerable cases. In fact this is not so; but there might be reasonable arguments that it is. Or, taking on board the hypothesis I floated about Yasmine, there might be arguments that it's so which even if flawed still strike Zachary as having some merit.

How does Zachary's epistemic standing with respect to P in the morning compare to his standing later in the day, on either of these variants? One answer, which doesn't satisfy me, is that his justification is in no way diminished. After all, he proved P. How much better reason could he have? This is what some hard-line anti-subjectivists would say.

As in the previous case, there are lots of details I've left unspecified. Different ways of filling them in may invite different judgments about how much Zachary should back away from his initial conclusion. For example, I wouldn't want to rule out the possibility of cases where Zachary can *just see* that his proof is correct, and the supposed experts massed against him must be wrong. But even

in that case, and certainly in many other cases compatible with the outlines I sketched above, I think, and I think most of us will agree, Zachary should be *somewhat* less confident in P, after his colleagues' testimony, than he was in the morning. Even if, as a matter of fact, he has proved P. If you like, you can even imagine *he proved that* he proved it. Still, the colleagues' testimony could be such as to make him rightly less sure.

Some are tempted to say here, "Well the colleagues' testimony at least gives him a reason to double-check his work." And I think this suggests an important insight, which we'll come back to later. But it isn't yet directly responsive to the question I'm working with. Perhaps Zachary does have reasons to double-check his work, or analyze his colleagues' motives or possible misunderstandings, or do many other things. Still, before doing any of that there is the question of what his epistemic status is right now towards P. The hard-line anti-subjectivist says that since Zachary arrived at P by a valid inference, he has the best kind of evidence possible and must be maximally justified. So even *before* or *without* double-checking his work, they'd say Zachary would be unreasonable to be any less confident. The hard-liner can say things to ameliorate this incredible claim. For example, they can allow that it may be reasonable for Zachary to doubt other claims in the neighborhood, such as the epistemic claim that his belief in P is justified. But the suggestion that first-order and epistemic claims can come apart like this brings in new controversies. (Myself I want to allow for it.) And we're still left with their intuitively unpalatable claim that Zachary should still always continue to be certain that P. Unless he, for example, forgot how the proof went.

This is one of the places where thinking about chunky attitudes can blind us. For there is a temptation to say, if Zachary *knew* that he proved that P, or perhaps even, if he just *knew* that P, on the basis of proving it, he shouldn't then back off when confronted by his colleagues. At least, not if he's retained the evidential basis on the basis of which he knew it. We'll come back to this in a moment. For now, though, let's stay with graded attitudes. I think the only credible position here is that, in some (or maybe all) cases of the sort I've described, it'd be reasonable for Zachary to be somewhat less confident. The hard-liner's position is one that no one should embrace willingly.

The subjectivist has an easy time accommodating the judgment that Zachary should be less confident. In the morning, Zachary's evidence consisted of his I've-proved-it phenomenology, which on its own does make P quite likely. But in the afternoon or evening, the memory of that experience has to share the field with the colleagues' testimony, and this combined body of evidence makes P less likely. At no point does Zachary's evidence itself *entail* that P; so it's possible for the evidence to be supplemented and the combined support for P come out less strong.

This thought cannot succeed in the form I've just stated it, because I'm imagining P to be a theorem, so (on most logics) it's entailed by every body of evidence. But there may be more promising ways to formulate the subjectivist's core idea

here.

But I'm dissatisfied with even the spirit of the subjectivist's approach. I agree we want to think of Zachary's epistemic situation in the morning as defeasible: that is, as being such that adding more evidence, without him forgetting anything, can leave him less justified in believing P. But we should resist doing so in a way that makes his epistemic situation no better than Yasmine's, or even Xavier's. As Zachary does in fact see, there's more to be said for P than just the way it makes him (and the others) feel.

So the position I like is one where: the justification one gets from proving P is nonetheless (i) defeasible, but (ii) is in some ways still superior to the justification one gets from falsely seeming to prove P, and (iii) each of these is better than the inductive baseline of how likely a proposition like P is given that you feel the way you do about it.

Anyone who admits point (i), that Zachary has reason later in the day to be somewhat less confident, admits that logically valid transitions needn't be epistemically conclusive, even when one is aware of their validity. (Being aware of their validity doesn't mean you're in a position to reasonably be *certain* they are correct.) And that is one of the main claims I want to stress.

(Perhaps Zachary has less justification to believe P *because* he has more justification to believe his belief in P is unjustified. If so, that's OK by me. Williamson has argued recently that knowing P is compatible with having arbitrarily strong evidence that you don't know P. So he is opposing any general principle to the effect that higher-order anti-justification suffices for a lack of a first-order epistemic property. I am sympathetic to that. Unlike many philosophers, I think it's possible to have fully justified beliefs that you're justified to degree u in believing P, and yet in fact be justified to some different degree v in believing P. However, even if we agree with Williamson that the higher-order anti-justification doesn't *suffice for* a lack of the first-order epistemic property, that still leaves room for the plausible claim that the higher-order epistemic claim has *some* negative first-order negative effect. It might still *contribute towards* your being less justified at the first order. And that is the proposal floated at the start of this paragraph. But whether Zachary's diminished justification for P comes via a higher-order route or not, what I want to stress is that it is there.)

6. Defeating belief and knowledge

In the second set of Zachary examples, I urged that we think of his justification for P being to some extent diminished by the evidence of his colleagues. Certain lines of resistance to this can seem natural if our attention is on chunky states. I want to identify them, and encourage ways of thinking about how the chunky states relate to graded states, that permit my claims about the defeasibility of Zachary's justification to stand.

Suppose Zachary starts out justifiably believing P, in response to a defeasible body of evidence. We then supply Zachary with defeating evidence, bit by bit, until eventually he is no longer justified in believing P. Focus on some intermediate point in this process. At this intermediate point, the net evidential support Zachary has for P has been somewhat weakened. Yet he may still be more justified in believing P than in retracting that belief. So it's not yet true that he's lost the chunky status of being justified in believing P. In practice, I notice that many theorists will here shy from saying Zachary's evidence has been "defeated" (or, if it's undermining evidence we're adding, that his evidence has been "undermined"). They hear these as a kind of success term. And since Zachary is still at this point justified in believing P, they judge that the defeating has been unsuccessful.

It may be that this is the most natural use of these words in English. However, we shouldn't think that Zachary at this intermediate point is on epistemically equal footing with Zachary at the start of the process. After all, we may just need to give intermediate-Zachary three new pieces of evidence against P to render him no longer justified. Whereas those three new pieces of evidence could fail to be enough to render starting-Zachary unjustified. So even though both Zacharys have the chunky status of "being justified in believing P," intermediate-Zachary's possession of this status is more precarious, due to the other defeating evidence we've slowly been prepping him with. So we should say that *all* of the defeating evidence here has *had some defeating effect*, even if sometimes that effect isn't yet strong enough to move Zachary from chunkily being justified to chunkily no longer being justified.

Knowledge may follow just the same pattern. Or it may introduce extra complications. Let's proceed carefully.

Minimally, most epistemologists will accept that it's possible to know P on the basis of a body of evidence that only defeasibly supports P. It's natural to suppose then, that the knowledge so formed will also be defeasible. But I have seen a number of theorists resist this step. They think that once someone knows P, additional evidence can't take the knowledge away. Now this dialectic gets complex because we may have to contend with contextualism or extra-evidential conditions for knowledge, and these theorists may allow that knowledge *can go away* as new evidence comes in. But they'll deny that the mechanisms that are making this happen count as "defeat."

A somewhat different strategy would hold that one can only acquire knowledge that P on the basis of a body of evidence strong enough to support justified belief in P, but that this knowledge can be retained even through changes that render one's evidence too weak to support fresh knowledge, for someone just then coming to believe P. I don't find that position at all persuasive, but I admit it is coherent.

Despite this possible resistance, though, I'm going to make the bold step of assuming that if knowledge can be had on the basis of defeasible evidence, then

knowledge is itself defeasible. That is, you could know P, but then go on to acquire more evidence whose net effect is to make you no longer in a position to know P.

I trust that most epistemologists will be disposed to go along with me about this. However, what I'm proposing doesn't sit so comfortably with some other things we say about knowledge. If knowledge is defeasible in this way, then presumably it should also be possible to give someone not defeating evidence but rather strengthening evidence. Or it should be possible to give them evidence which immunizes them against defeating evidence they might acquire but haven't yet acquired. (We pre-emptively defeat some of those defeaters.) All of this suggests that just as knowledge is defeasible, so too is it improvable. But that's a thought I know many epistemologists will resist. My attitude is that the substance of the thought is correct, but there are probably semantic constraints on "knowledge" that we haven't yet theoretically mastered that make it tricky to figure out how to express the thought. Here is a try: when you know P, it is possible for you get into a stronger evidential position with respect to P. Someone in that stronger position would not be on equal epistemic footing with their earlier self, in the way that starting-Zachary and intermediate-Zachary in our preceding discussion weren't on equal epistemic footing—even though both subjects may know P, and both Zacharys were justified in believing P. Despite all this, for some reason we don't count the stronger epistemic position as "having more knowledge," or "knowing better"; and we don't count it as being in a position better than knowing, either.

But I may be wrong. Perhaps knowledge, even though it can be had on the basis of defeasible evidence, brings with it a kind of "maxing out," so that additional evidence can't effect any justificatory improvement, not even improvements described in terms other than knowledge. I don't find that at all plausible, but again, I admit it is coherent. I'm not sure I can understand the motives of someone who wants to say knowledge is not improvable. But I'll observe it's at least *coherent* for them to go on to say, in the other direction, that knowledge is after all defeasible. What one says about improvability doesn't logically mandate what one must say about defeasibility. And it's the claim that knowledge is defeasible that I want to proceed with.

I point these options out because they may tempt some readers to resist the idea that knowledge is defeasible. I prefer that temptation to be out in the open where we can confront it directly.

In any event, if I'm right that knowledge is defeasible, then we can expect it to behave in something like the way "justified belief" did. That is, Zachary starts out knowing P. We feed him some defeating evidence, but haven't yet rendered him *enough* less justified that he no longer knows P. In chunky terms, we may then shy from saying his knowledge has yet been defeated. But, as in the earlier example, we should be ready to grant that Zachary is now in a more precarious epistemic position. It will now take *less* defeating evidence to render him a non-knower than it would have taken if we hadn't interacted with him.

So the evidence we've fed Zachary *had some defeating effect*, even if it wasn't yet strong enough to move him from the chunky epistemic state of knowing to the chunky state of not-knowing.

If we keep all of that in view, I hope it will soften one line of resistance to my claim that:

- (a) Though Zachary may believe or know P on the basis of proving it, his epistemic position towards P will be weakened by his colleagues' testimony that his performance was flawed, or that the correct logic does not validate P.

For you may just be having a hard time imagining that the defeating evidence Zachary gets could be strong enough to render him no longer chunkily justified in believing P, or no longer in a position to know P. That is, you may be thinking:

- (b) Despite his colleagues' testimony, Zachary would still be justified in believing P, and still be in a position to know it.

Well, perhaps so. I'm not insisting otherwise. I hope that you'll now be in a better position to see that the kind of defeating effect posited by (a) can exist even if it's not strong enough to falsify (b). On the other hand, if you are now willing to grant me (a) but still want to hold onto (b), you have to face the question why it's not true that:

- (c) If Zachary can get a little bit of defeating evidence, he can also get a lot.

That claim is not inevitable, but it is *prima facie* plausible. And it doesn't seem like all of (a)-(c) could be true. I'm not going to try to sort out for you whether you should give up (b) or satisfy yourself that (c) isn't true. My aim is only to help you see that a commitment to (b) needn't prevent you from acknowledging (a).

7. Different axes of epistemic assessment

Let's turn now to the first set of Zachary examples. Recall that Zachary believes P on the basis of a genuine proof, Yasmine on the basis of a compelling but flawed proof, and Xavier on the basis of some opaque source like a crystal ball. I tentatively proposed that even if they all have the same track-record, we might think Xavier has reason to be confident to degree x in P, Yasmine to degree y , and Zachary to degree z , where $0.5 < x < y < z$. I also indicated some reservations with that proposal. I want to develop those reservations, and suggest a different

model for thinking about how these subjects are epistemically related to each other.

It's problematic to understand reasonable credences in terms of reasonable betting dispositions, for a variety of reasons. It's especially problematic to do so when thinking about reasonable credences towards non-contingent claims, as we are here. I think many of us will have well-entrenched intuitions that Zachary is at least more justified in believing P than Xavier is. However, even that much is hard to secure on existing formal models. What we write in our textbooks says that where P is a tautology, all the subjects can only reasonably have maximal credence in it. There are moves we can make to resist that. But it's not easy to avoid swinging to the other extreme, and saying that Zachary has no more justification for P than anyone else whose source has (and is reasonably believed to have) as good a track record. Even though Zachary seems to—and does!—discern the logical necessity of P, and the other's grounds for believing P offer no such insight.

I don't know whether we can, or should, talk ourselves into the thought that it would, after all, be reasonable for Zachary to accept higher odds on P than Xavier. (What would count as settling the bet, anyway? Hasn't Zachary already settled it, because he proved it? Or has he failed to settle it, because his justification for believing P is still defeasible? But then, could the bet ever be settled?)

But for the moment, let's just use that model to think about the relative standings of Zachary and Yasmine. Even if we do talk ourselves into the thought that Zachary would be reasonable in accepting higher odds on P than Xavier, do we also want to say it'd be reasonable for him to accept higher odds than Yasmine?

I find this extraordinarily difficult to bring myself to believe. Whatever advantages Zachary may have over Yasmine, I don't think they come down to its being reasonable for him to be more confident in P than she is, or for him to accept higher odds against P than she should. Both should perhaps reserve some credence for the possibility that they've messed up, but as I presented the case, neither has any more independent reason than the other to think they're going wrong in this case. (And nor do they disagree! They both think that P.) So how should Yasmine justify to herself being less confident in her result than Zachary should be? Why should she think she's more likely to have reached that result erroneously?

Now a philosopher with a strongly externalist enough a stomach may be able to reconcile themselves to this. But I think many will share my reluctance, even if they don't keep to an entirely internalist diet. But if we do deny that it'd be reasonable for Jasmine to have less confidence in P than Zachary, what then becomes of the thought that in some sense Zachary's position is epistemically superior?

The key is the qualifier "in some sense." Perhaps how much confidence your position justifies in P is not the sole measure of epistemic strength.

Of course, there is also how much confidence your position justifies in *other* propositions. Perhaps the difference between Zachary and Yasmine could be explained in that way: maybe he has more justification to believe he's *proved* P than she does. Or maybe he's more justified in believing that *he is* justified in believing P to the degree he is, than she is in believing the corresponding claim about herself. Or so on. I don't know how satisfying those strategies will turn out to be. I want to explore other ways of implementing the root thought.

Another way to go would say that Yasmine's justification is less *resilient* than Zachary's: that is, she's vulnerable to kinds of defeating evidence that he isn't, such as evidence that points out her error. But on the other hand, he may be vulnerable to defeating evidence she isn't: he may have a genetic condition she lacks which makes him quicker to succumb to hypoxia, and if so, evidence that suggested there's too little oxygen present would defeat his justification for P more than hers. So it's tricky to say in what sense her vulnerability to defeating evidence outweighs his. I agree with the intuition that it does, but we can't just count up pieces of evidence that would threaten the one of them and not the other.

It may be that we can find some other familiar epistemic properties in which Zachary's advantage over Yasmine consists. Even if we can't, I would be willing to accept it as a normative primitive that his justification is *of higher quality* than hers, even if they both mandate (or invite, or permit) the same degrees of confidence. This is not an unprecedented proposal: some philosophers working on testimony have proposed, analogously, that learning P from a trusted trustworthy testifier puts one in a position that's epistemically superior to that of someone who just eavesdrops on the same conversation—but this superiority shouldn't be understood in terms of mandating more confidence. (This is Miranda Fricker's view. I don't know who else. Some of the philosophers who think testimony is special do express that in terms of mandating more confidence, but I doubt that's the most hospitable frame for their other views.) Another analogy is the sense many have that our justification for a mathematical claim would be somehow better if we had proved it ourself, than if we established it by computer-aided enumeration. Though we could arrange the reliability of these methods so that we should hesitate to think the first called for more confidence than the second.

One way to avoid taking this superiority as primitive might be to say that Yasmine ought to double-check her reasoning more vigorously than Zachary—again, without this yet implying that she should now be less confident. I'm attracted to the idea that epistemic norms about when you should seek more evidence, or corroboration, or how vigorously you should review your reasoning, aren't strictly correlated with how confident it's reasonable for you to be at the moment, before those things are done. So I think it's quite plausible that Yasmine and Zachary are now reasonably confident to the same degree in believing P, but that some of these other, somewhat orthogonal or oblique norms weigh on them differently. And that might explain our sense that Zachary's epistemic

position is better than Yasmine's.

And in fact, such axes as these might also be relevant to the presence or absence of some chunky attitudes. It's not implausible that whether you count as "believing" or as "still making up your mind" is tied up with your dispositions to seek more evidence, corroboration, and so on. Hence, perhaps whether you count as "justified in believing" is tied up with how much more evidence, corroboration, and so on, you *ought* to be seeking. And so we might persuade ourselves that, though Yasmine and Zachary are justified in having the same degrees of confidence, Zachary but not Yasmine is justified in settling on this verdict, and chunkily believing P. Whereas for Yasmine, belief might be premature, because she has fulfilled fewer of her intellectual obligations in getting to this reasonable confidence than Zachary has.

I don't want myself to undertake a defense of these proposals. But they do sound viable to me, and they help confirm my impression that it could make sense to count Zachary as being *in some sense* epistemically better off than Yasmine, even if it's not the sense of being reasonably more confident.

That's even further confirmed when we observe that we'd count Zachary as in a position to know P, but we wouldn't so count Yasmine. (But I really don't understand *why* we have this judgment about Yasmine. I'd rather work my way towards it than use it for leverage.)

8. What supports what?

My final group of arguments will have a different focus. Logic concerns the relation between sentences or propositions (if propositions have sufficient structure). And many epistemologists believe that justification or evidential support is also a relation between propositions. However, this latter view is incorrect, and if I can discourage it, it may help make clearer how non-straightforward the connections are between valid inferences and justified inferences.

I'll begin by observing that the support relation has a left-hand side and a right-hand side. I will mostly be discussing the left-hand arguments. But we shouldn't think the right-hand arguments are propositions either. Some epistemic positions support withholding judgment, rather than any proposition. And this can't in general be explained as a matter of one part of your evidence supporting *one proposition* and another part of your evidence supporting another. Your epistemic position might support withholding judgment about P because it supports believing that *you're not competently placed to judge whether P*. Your position here does support believing *something*, but by doing so it supports an attitude towards P that's not itself a proposition, nor a belief in any proposition.

And support for withholding can wax and wane: it may be that I start out justified in believing neither P nor not-P. Then later I gather more evidence

that supports withholding judgment about P even more strongly. For example, I may know now, which I didn't know before, that P is not just unknown but unknowable.

One might attempt here the view that your position still supports a proposition about P, it's just not the proposition P and not the proposition not-P. Instead, it's the proposition *Maybe P and maybe not P*. We'll come back to this proposal in a bit. For the moment, I'll just remark that it seems more promising to understand the "belief" that *Maybe P and maybe not P* in terms of an explanatorily more basic attitude of suspended judgment on P, than it seems to understand the latter in terms of belief in an exotic proposition.

There are other epistemic responses besides belief and suspended judgment. Earlier in this paper, we considered responses like not-making-up-your-mind-now, or closing-some-inquiry, or seeking corroboration. We could add responses like doubting or suspecting, which aren't obviously a matter of disbelieving or having some given graded confidence. I am not sure whether these kinds of epistemic responses are mandated or supported by our evidence. If they are, they may not be "supported" in the same way that beliefs and withholdings and graded confidence are. But I'm not sure about any of that. Certainly we're willing to say that some evidence "calls for doubt," or that other evidence settles a question. How these expressions should be theoretically interpreted is not obvious. But it is some evidence that we think of doubting and settling as being among the right-hand arguments of the support relation.

The rest of our discussion will be devoted to the other side of the support-relation. What are the possible providers of support?

Here are five arguments to think that it's wrong to assign this role to propositions.

- Sometimes *your lack of evidence* supports withholding judgment. In saying that, I'm not saying that the proposition that you lack evidence is supporting anything; at least, I'm not saying that the proposition that you lack evidence is playing the kind of role that the proposition that past emeralds have been green plays towards the conclusion that the next emerald will be green. That is, I'm not saying "I lack evidence" is anything like a premise which you should take to be propping up some conclusion. (There is here the additional difficulty that it's not clear what the conclusion should be.) There *could be* cases of that sort, but I take it that's not what's generally the case when your lack of evidence supports withholding judgment. Generally it's some global *pattern in* your premises, rather than *some additional premise* in the pattern, that is doing the epistemic pushing.
- Similarly, some may say there are cases where your *tiredness or incapacity* supports withholding judgment. This has to be handled more carefully. For here it's more plausible that your tiredness or incapacity does

standardly engage with your reasoning as a premise. And some theorists (perhaps including me) will say that's the only way for these things to normatively bear on what you're justified in believing. But that view is not inevitable. It's also quite intelligible to hold that the bare facts about how tired you are, whether they're available to you as premise or not, affect how justified you are. Those who embrace this view will say that your tiredness itself supports a doxastic attitude, in something like the way that *your belief* about past green emeralds supports a different doxastic attitude. Their view *won't* be that your tiredness plays the role also played by past-greenness. Rather, it plays the role also played by your belief about past greenness. Or perhaps the role played by your justification for that belief. So in the more familiar emerald case, it isn't the proposition about past-greenness on the left-hand side of the support relation, but rather some attitude or epistemic fact about you.

- *Withholding* about one thing can support other attitudes towards other things. For example, suppose I'm in a situation where I know that exactly one of seeing or hearing is reliable, but it's not yet settled which. Seeing tells me P, and hearing tells me Q. Suppose I've gotten into a position of believing it's probably my visual experiences which are genuinely perceptual. But then I change my mind, and decide to suspend judgment about P. Given what else I know here, won't that affect what doxastic positions I can coherently or reasonably have about Q? Won't my suspension of judgment about P put normative pressure on me to raise my confidence in Q, that is, to go from disbelief in Q to suspended judgment?

One might attempt here the view that it's not my suspension of judgment that is doing the normative pushing, but rather the evidence that persuades me to suspend judgment. But what if there is no such evidence? What if I just change my mind on a whim? Or because I see some old evidence in a new light? Don't these mere changes of mind, even if not rationally generated, impose some coherentist pressure on me to change other parts of my total opinion?

There are thorny issues here. My own view is that it is important to distinguish the supporting relation that is exemplified by actual evidence or justified belief, on the one hand, and the "coherentist" normative pressure that it is exemplified by mere beliefs or changes of mind, on the other (regardless of the latter's epistemic merit). So I am sympathetic to those who say the case where you suspend judgment not in response to evidence is importantly different from the case where you do so in response to evidence. At the same time, I think we need to understand these cases in a way that's theoretically connected. And I think in the end the mere attitudes end up having *a* kind of epistemic role, though it's not the same as the role had by evidence or justified attitudes. But this takes us deep into issues that I've already said are thorny. So let's leave it for another day.

- There are several classes of assertions whose semantics might best be understood non- or extra-propositionally. This is all quite controversial. But one large tradition in linguistics over the past quarter century has to been to treat assertions like:

1. Sarah might be in the cupboard.

or:

1. If Sarah is outside, then the cat is in the cupboard.

or:

1. If Sarah is outside, she's probably riding her bike.

as fundamentally different from assertions like:

1. Sarah is in the cupboard.

If we only had to deal with the latter, we might get by with identifying the assertion's content with which possibilities it leaves uneliminated. But on the proposed accounts of the other assertions, this does not suffice. For example, the semantics proposed for the first says: in a situation where some Sarah-in-the-cupboard possibilities are currently alive, the assertion does not eliminate (and neither does it add) any new possibilities. If on the other hand, there are no Sarah-in-the-cupboard possibilities currently alive, this assertion eliminates every possibility. That is, in that case no possibilities are compatible with what's been asserted. This is the "dynamic semantics" treatment of epistemic modals championed by Frank Veltman and others. There are also important proposals in this tradition for conditionals, anaphora, presupposition, imperatives, and other constructions.

Nate Charlow has been developing an account of expressive utterances that's modeled on some of those accounts of imperatives. Without getting into the details, I'll just observe that there are at least superficial similarities between the phenomena just mentioned and a Gibbard-style account of some expressive language. Gibbard's account of moral language is of course quite controversial, but some might be better-disposed towards a structurally similar account of exclamative expressions. For instance, there seems to be some difference in meaning between:

1. Tom gave me the damn book.
2. Tom gave me the book.

Yet it's not clear this difference is well-understood in terms of different truth conditions. Rather, it seems like "damn" has an effect on the meaning of the assertion that is essentially a side-effect, rather than part of the

normal business of assertorically ruling out possibilities. Charlow's work suggests that this similarity is more than superficial, that these examples are theoretically unified with the others.

There is a large literature on all of this. I am only trying to gesture at it, I hope in a way that will still be usable by those to whom this is less familiar. The bottom line is: there are a variety of linguistic pressures that motivate some theorists to posit meanings for these assertions that can't be reduced to the kinds of truth-conditional content we're familiar with for:

1. Sarah is in the cupboard.

These accounts are not universally accepted. At the same time, they aren't fringe proposals either. They are mainstream theories that have proven extremely fruitful. So they cannot reasonably be ignored, either.

Well, what of it?

Where I am going with this is that, even if these assertions don't have truth-conditional meaning of the sort we're familiar with, yet they don't seem to be epistemically inert either. For example, the first chain of reasoning seems rationally more intelligible to me than the second:

1. Tom didn't give me his damn book.
Sarah didn't gave me her damn candle.
So I don't have a damn book or a damn candle.
2. Tom didn't give me his damn book.
Sarah didn't give me her candle.
So I don't have a book or a damn candle.

I don't claim there is *no* situation where the second discourse could make sense. Only that it would make different sense than the first. The pattern of the "damn"s does seem to affect what thought transition is being envisaged.

Similarly, I can easily imagine being in a position to reasonably make this inference.

1. Tom has the same evidence I have concerning Sarah's whereabouts.
Tom says Sarah must be outside.
So if Sarah might be in the cupboard, then Tom is lying (or less clever than he thinks, etc).

The epistemic modals are playing an essential role in the inference I'm here envisaging. This is all far from decisive, but it does suggest that whatever semantic exotica we need to account for the meanings of these assertions will also be capable of getting into the support relation.

Now if the relata of the support relation are attitudes, this won't be that difficult to manage. One of your beliefs might be:

1. Sarah might be in the cupboard.

And we can then go on to ask, what other beliefs does that initial belief support? But we won't then be asking about what truth-conditions are supported by *the truth-conditions of* the initial belief, because on these proposals, the initial belief is *not* rightly understood in terms of what makes it true, or what possibilities it does and doesn't rule out.

If on the other hand the relata of the support relation are restricted to propositions as familiarly understood, then it will be more difficult to understand what's going on when the belief that Sarah might be in the cupboard persuades you that Tom is lying.

There are many places here that one can resist the line of thought I'm sketching. You might in the first place think no exotic semantics is needed for the kinds of language I've mentioned. You would have some good company in thinking so, but you would also have many formidable opponents. Or you might allow that exotic things are going on with the language, but resist my suggestion that it can affect the epistemology. I'll be interested to see how that turns out.

Or, finally, you might accept everything I've said here, but deny that it bears on the connection between logic and epistemology. You might say, instead, this all just goes to show we need a different (a structurally richer) notion of proposition than we started with. If:

1. Sarah might be in the cupboard.

is something you can believe, then we should count it as being a proposition by stipulation. The exotic semantics only teaches us better what propositions are, not that propositions aren't the relata of the support relation.

I sympathize with this last proposal. "Proposition" is a term of art that's been conceptually associated with many different job roles. As certain parts of philosophy evolve, it's becoming less obvious that the same entity is suited to play all those roles simultaneously. For presentational purposes, I've been treating "propositions" as correlated with truth-conditions. We didn't get into details, but I was minimally assuming that propositions stood many-one to truth-conditions. Yet the semantic proposals I mentioned are denying that the contents of assertions can be understood in those terms. It is a step, but a natural step, to think then the contents of beliefs cannot be understood in those terms either. I acknowledge another way to go here is just to designate the contents of assertion and belief, whatever they turn out to be, the propositions.

What we should care about for our current discussion is not what gets to be called "proposition," but rather what are the relata of logical entailments, on the one hand, and the relata of the support relation, on the other. Some of the other arguments I gave were pushing for the relata of the support

relation not being belief contents at all, no matter how exotic. (Though recall the idea, when we discussed withholding, that perhaps withholding judgment about P is the same as believing *Maybe P and maybe not P*. I won't argue against that idea, but just note its connection to the present topic.) On its own, the argument I've developing here doesn't purport to do that. On its own, it would at best show that the relata of the support relation have to have at least the structure had by belief contents. I was pointing out that's more structure than is had by truth-conditions. (And though we didn't go into the details, these semantic proposals require it to be more structure than is had by a strict refinement of truth-conditions.)

Now, if logical entailment is also a relation between arguments with that much structure, too, then the force of my present argument would be blunted. And I admit this is possibly how this dialectic should play out.

Most of the work I'm aware of about epistemic modals, conditionals, exclaimatives, and so on has dealt with their semantics, not their proof theory. I know of only a segment of the literature that deals with their logic, and it's not even clear that it is dealing with a logic rather than a programming language. (Not that that difference is so easy to articulate.) In any case, there could easily be logical work out there that I don't know of, or it may be just around the corner. There is no formal obstacle to our extending logic to operate on the kinds of richer semantic structures that are posited to be the meanings of epistemic modals and so on.

Of course, it's not enough for our purposes that there be *some* formal theory of this. As we saw in section 4, we have *too many* logics. But it could well be that the *correct* logic (whatever we end up agreeing *that* means) is one that operates on these richer arguments. And in that case, I will concede that the present line of reasoning does not advance my case that logic and justification are less straightforwardly connected than is nowadays assumed. I will in that case stand by the other arguments, instead.

Regardless of whether it advances my case, though, I think these issues are intrinsically important and interesting for epistemology. If we do get to the dialectical situation I've just described, we will for instance need to think differently about the relations between belief, justification, and truth, than many nowadays do. This is all worth discussing in its own right, on another occasion.

- I have one last argument to offer against the left-hand relata of the support relation being propositions. Actually, it's a more an argument template. The template goes like this:

So-and-so's theory of justification says that X supports doxastic response such-and-such.

X either has no content, or its content doesn't support the content of that doxastic response.

We're not yet in a position, as theorists, to decisively exclude so-and-so's theory.

I will give four examples.

- (a) Alex Byrne says that believing P supports the belief that you believe P. — Well, it's not *entirely* obvious he says that. What he explicitly says is that the transition from the belief P to the belief that you believe P is automatically rational. I suppose it's not straightforward what is the relation between some transition's being rational and the support relation being exemplified. But we've waded through complexity enough, so let's just imagine Byrne will say this about the support relation. Yet does the proposition P support the conclusion that I believe P? There are some propositions for which that clearly does not seem correct. For instance, the proposition that I unknowingly have a "Kick me" sign on my back doesn't support the conclusion that I believe I unknowingly have a "Kick me" sign on my back. Neither does the proposition that Socrates' blood type was O-negative seem, on the face of it, to support the proposition that I believe it was. At least, if *you* were to discover what his blood type was, it would be rash of you to believe I already believed it. These don't by themselves have to refute Byrne's view. But they suggest that, if he's going to talk about support at all, the kind of epistemic relation he's proposing is better thought of in terms of *beliefs* (or other attitudes) doing the supporting, rather than propositions. *My belief* that Socrates' blood type was O-negative seems a better candidate for supporting the conclusion that I have that belief. (And it doesn't seem like I *could* believe I unknowingly have a "Kick me" sign on my back. Unless I think I'm in a Gettier case or something.)
- (b) Others hold views like Byrne's about cogito judgments. They say that what justifies the beliefs that I exist, and that I am thinking, is not any inference, nor either some occurrent phenomenology, but rather the self-verifying nature of the cogito judgments. I have argued against these theories elsewhere. But I submit them now anyway, as exhibits of epistemologies we may not be in a position to decisively exclude, that seem ill-suited to be expressed in terms of some propositions supporting other propositions. For in these cases there presumably are no left-hand supporting propositions. But it doesn't seem like the proposition that I am thinking is supported by the null set of premises. One objection: *You* aren't necessarily justified in believing I'm thinking. It may be responded: you are justified in believing the same "de se" proposition about yourself. Ok, let's not pursue that. A second objection: what if I am *not* thinking? Instead, I'm in a state of perfect meditation. Do I then have justificatory support for the de se proposition that I *am* thinking? To be sure, I wouldn't go wrong if I were to form that thought, for I'd thereby make it true. But

it seems unhappy to understand this in terms of my already having justificatory support for its content. Rather, the kind of epistemic relation being envisaged here seems better thought of in terms of *the belief* that I am thinking being supported by the null set of premises.

- (c) Another view like Byrne's would say that the mere experience of imagining a table should justify the belief that I imagine a table. That is, if I have that occurrent mental imagery I am automatically justified in believing I do, in response to or on the basis of that imagery. Perhaps this is because the experience is essentially a conscious one; or perhaps it's because imagining is essentially a mental activity of mine. I think we can leave those details unspecified. Here too it is hard to explain the posited support relation in terms of propositions doing the supporting. What should the supporter be, and what supported? It's not the propositional content of my mental imagery that's doing the supporting, for I don't think mental imagery *has* propositional content. It's structurally analogous to a noun phrase, rather than to a sentence. But even if you think the imagery *does* have propositional content, arguably that would be some proposition about a table, not about my mental relation to it. (I say this is arguable, not that it's inevitable. I think it is the position most who chose this fork would favor.) And why should a proposition about the table, on its own, support the claim that I'm imagining the table?

If on the other hand, we say it's not any *content of* the image that's doing the supporting, but rather *the proposition that* I'm having that image, then we'd have one and the same proposition on both sides of the support relation. Some will think that's already intolerable. I'm not sure it is intolerable. But we can ask: if the proposition that I'm imagining a table supports itself, and thereby justifies me in believing that I'm imagining a table, why shouldn't the proposition that I'm imagining a flying pig (which, let us suppose, I am not) also support itself? That is, what stops me from being justified in believing, falsely, that I'm imagining a flying pig? It seems we'll have to say it's only when the proposition that I'm imagining a table *is true* that it supports itself. But then what is the advantage of saying that, over saying that what's doing the supporting is not the proposition but rather *the fact* of my imagining the table? Or alternatively, the mental act of so imagining? Or the mental imagery? I don't claim these latter alternatives are inevitable. The theory that it's *the proposition* that's doing the supporting, but that it only does so when it's true, is intelligible to me, and I understand the technical difference between it and the other proposals. But it's hard to understand what should motivate it.

- (d) David Christensen thinks believing I took a bad-at-math pill, or am in conditions conducive to hypoxia, which messes up a subject's reasoning in ways the subject then lacks the capacity to detect, can

support being less confident in some result P that I've just deduced. But P might here be a claim of pure mathematics. Do we think that the contingent propositions about what drugs I've ingested or how much oxygen I'm getting epistemically bear on what's true mathematically? I agree we should think that *believing* those things about yourself should affect how reasonable it is for you to *believe* the mathematics, or what graded confidence it's reasonable for you to have. But that thought doesn't seem so happy when formulated as a claim about what *propositions* support what other propositions. As with the Byrne proposal, if *you* learn the propositions about my incapacities, it doesn't seem like *you* should think P is less likely to be true. (I assume your current confidence in P is not based on my say-so.) Similarly, if I learn that I will in the future suffer those incapacities, it doesn't seem I should think P will in the future be less likely to be true. Or if I hypothetically suppose I suffer those incapacities, without having any evidence that I do, it doesn't seem that I should infer that in that hypothetical scenario, P is less likely to be true. So here again, it seems like the best way to understand the view is in terms of it being *my beliefs* (or other occurrent, non-hypothetical doxastic attitudes) that are doing the supporting, rather than propositions.

Anyway, there are some examples. I'm not especially sympathetic to the first two proposals; I'm middling on the third; and I am very sympathetic to the fourth. But even if you indisposed towards all of these proposals, I hope you'll at least agree we're not yet in a position as theorists to decisively rule them out. At least, the field at large isn't. And each of them seems to invite us to think of the epistemic support relation in terms of *what facts about a subject*, or *which of the subject's attitudes or experiences* or so on, support what doxastic responses. Rather than in terms of *believed contents* doing all the supporting.

I haven't been making a trivial point about theoretical bookkeeping. It's not as though we could choose to say it's either the proposition that the past emeralds were green, or my belief that the past emeralds were green, that supports a conclusion about future emeralds—and having made either choice, trivially translate into the other terminology.

No: the objections I've been developing testify that the view that the propositional contents of my premise beliefs are the supporters is *too structurally impoverished* to permit many of the varieties of epistemic support we should want to allow.

On the other hand, logical entailment is presumably a relation among belief contents—even if they are the richer contents some have proposed for epistemic modals and so on, rather than the merely truth-conditional contents we're more familiar with.

Of course we should expect there to be interesting relations between these. For example, nothing I've said so far excludes the hypothesis that when P entails Q, some starting position associated with P—perhaps the belief that P, or having justification for P—should give you *prima facie* justification to believe Q, justification which might however be defeated by other evidence, such as the news that you've taken a bad-at-Q pill, or testimony from your colleagues that P doesn't entail Q on the "correct" logic, or so on. I'm not going to defend or reject that hypothesis. I'm just pointing out that more complex connections between logic and justification than I've been arguing against here may well exist.

But I hope to have at least shaken the simple idea that support is a relation between propositions, and that entailment—another relation between propositions—is one species of, and so suffices for, support.