

The ordinary language argument against skepticism—pragmatized

Sinan Dogramaci¹

Published online: 18 December 2018
© Springer Nature B.V. 2018

Abstract I develop a new version of the ordinary language response to skepticism. My version is based on premises about the practical functions served by our epistemic words. I end by exploring how my argument against skepticism is interestingly non-circular and philosophically valuable.

Keywords Skepticism · Ordinary language · Function · Ambitious anti-skeptical project

1 Setting up the target views

1.1 Skepticism, anti-skepticism, and evidence neutrality as semantic theses

Skepticism is the view that we know little to nothing. Anti-skepticism is the view that we have plenty of knowledge. We can reformulate these as semantic theses by taking a step of semantic ascent. Skepticism then becomes the view that our word ‘knowledge’ correctly applies to few or none of our beliefs; that is, it has few or none of our beliefs in its extension. Anti-skepticism becomes the view that ‘knowledge’ has lots of beliefs of ours in its extension.

Let’s also reformulate an important premise in the skeptic’s argument for skepticism. The skeptic makes their argument not by arguing that our beliefs are flat out false, but by arguing that they aren’t justified. In other words, the evidence we possess, the skeptic claims, doesn’t adequately support the propositions we believe;

✉ Sinan Dogramaci

¹ Department of Philosophy, The University of Texas at Austin, Austin, USA

or in yet other words, our evidence doesn't make the things we believe very likely to be true. Different arguments for skepticism share a key premise of *evidence neutrality*. The premise says that all the evidence we have, or could ever obtain, lends no more support to our ordinary beliefs than it lends to various skeptical hypotheses that compete with our ordinary beliefs, such as that we're brains in a vat, that other people are mindless zombies, or that the laws of nature will suddenly break down.¹ The skeptic concludes that any evidence we have or could ever get will leave our ordinary beliefs at least as likely false as true. Since knowledge requires justification, indeed requires that what we believe is at least very likely true if not surely true, the conclusion that we can know little to nothing also follows.

Semantically reformulated, the skeptic's evidence neutrality premise becomes a premise about the extension of a relation, the *x-is-good-evidence-for-y* relation. The extension here consists of ordered pairs. The first member, *x*, is a body of evidence. (When we think of *x* as the evidence a person possesses, we'll understand this to be the person's *total* evidence. The reason: if only a part of a person's evidence makes a proposition likely, it might still be false that the person has justification to believe that proposition.) The second member, *y*, is a proposition, a hypothesis whose likelihood is in question. To make the vague notion of "good evidence" more specific, we can focus in particular on the relation where the total evidence, *x*, makes the hypothesis proposition, *y*, more likely than not. Then evidence neutrality, the key premise of the skeptic's argument, says this: the *x-is-good-evidence-for-y* relation does not contain in it pairings of (*x*) evidence that we actually or could possess paired with (*y*) very many, if any, of the ordinary propositions we believe.

The anti-skeptic, who denies evidence neutrality, then must endorse the following semantic reformulation of the premise's denial. The anti-skeptic endorses this: the words that we use (interchangeably) to express epistemic evaluations, 'justifies', 'supports', 'makes likely', and 'is evidence favoring', have an anti-skeptical extension, that is, they include in their extension lots of ordered pairings of (*x*) evidence we do or could possess, paired together with (*y*) ordinary propositions we believe.

1.2 Immediate benefits of giving semantic formulations

What use is it to entertain all these semantic formulations of skepticism, anti-skepticism, and evidence neutrality? I want to suggest that, even before we consider any arguments, just by turning our attention to the semantic formulations, we've made it easier to wholeheartedly embrace anti-skepticism. The point I want to make right now is a psychological point. I think it's simply a lot easier, psychologically, to appreciate *how* anti-skepticism could be true when we formulate it semantically. All it takes for anti-skepticism to be true is for our epistemic terms to have picked up anti-skeptical extensions. For example, 'supports' just needs to contain in its extension pairings of (*x*) evidence we do, or could, possess and (*y*) lots of ordinary

¹ See, e.g. the discussions by Pryor (2000) and White (2006). On the "could ever obtain" aspect of the argument, see esp. Pryor's pp. 528–531 and material he builds on in Stroud (1984).

propositions, like that I have hands and that the sun will rise tomorrow. Our other epistemic terms, including ‘evidence’, also need the appropriate anti-skeptical extensions. Sure, maybe it could be that we’ve ended up using words with skeptical extensions, but, equally surely it would seem, it could be that we are, and have been all along, using our words with anti-skeptical extensions. Why not? Can’t *some* words have anti-skeptical extensions, and why can’t *our* words have them? The skeptic does aim to argue that their view is necessarily true, in a metaphysical sense. The skeptic argues that, for creatures like us, no metaphysically *possible* evidence could justify our ordinary beliefs. But that necessity depends on it being *actually* true that our words have skeptical extensions. And that alleged actuality is made especially questionable when we think about the possibilities where our words have anti-skeptical extensions.

So, the mere existence of anti-skeptical extensions for our epistemic terms is useful to know about. But I haven’t provided a detailed example of such an extension yet. Is there a worry that such extensions don’t even exist? I don’t think so. The worry would have to be that there is no consistent characterization of a candidate anti-skeptical extension for our epistemic terms; in figurative language, such extensions aren’t even stocked on the shelf in Plato’s heaven. This sort of worry certainly has force for other philosophically contested notions. The deep challenge for ‘true’, and the concept of truth, in light of the Liar paradox, is that attempts to characterize their extension run into inconsistency. It thus takes hard work to even make it plausible that there’s a coherent candidate extension for ‘true’. But our epistemic terms don’t face such a challenge. That’s because the traditional skeptical arguments are not arguments that there is no possible candidate anti-skeptical extension for ‘justified’ or ‘knows’. The skeptics didn’t argue that there is no abstract object in Plato’s heaven that both is anti-skeptical and could remotely reasonably turn out to be the extension of ‘knows’. And we do have good examples of reasonable such candidate extensions. The examples are the extensions that come with existing anti-skeptical theories, like dogmatism, neo-rationalism, and the anti-closure view.² When I call these extensions “candidates”, I mean they look like extensions one of which reasonably might turn out to be our epistemic terms’ extension, given our actual usage. It’s reasonable to think philosophical examination and reflection might ultimately lead us, maybe after we’ve achieved wide reflective equilibrium, to conclude, for one of these, that it’s the actual extension. In that sense, I think we have good candidate anti-skeptical extensions available on the shelf, in Plato’s heaven, for our epistemic terms. And this is my optimistic starting point from which I want to argue that one or the other of these is the *actual* extension.

² See Pryor (2000) for dogmatism; see Wright (2004) and White (2006) for neo-rationalism; see Dretske (1970) for anti-closure.

2 Arguing for anti-skepticism via the pragmatics of language

So now let's try to argue for anti-skepticism. We'll take the semantic formulation as our targeted conclusion. We'll argue that our epistemic terms, 'knows' and 'supports' and the rest, actually have anti-skeptical extensions.

I'm following in the tradition of the so-called ordinary language school, whose anti-skeptical proponents included Malcolm, Austin, and Strawson.³ I'll be pursuing an argumentative strategy of my own, so I won't examine their particular views and arguments (or lack of arguments, actually) (and I'll relegate my response to an unjustly influential objection from Salmon to a footnote.⁴).

My proposed strategy is to begin with premises about the pragmatics of our epistemic terms, and using these premises make an argument for anti-skepticism, semantically formulated, as our conclusion. These premises I'll borrow from other work. The main step forward I want to take and defend here is to propose a pragmatics-to-semantics link, and to show how, given that link as a plausible further premise, anti-skepticism follows. But first, let me state those starting premises about the pragmatics of our epistemic terms.

2.1 Pragmatics of our epistemic terms

I use 'pragmatics' to mean views about what *function* our epistemic terms serve, what point or purpose they serve. I'm not using 'pragmatics', as many philosophers of language do, to mean views about how words can communicate information that differs from their semantic content. My use is more in the spirit of pragmatism, the philosophical tradition.⁵

So then, what function, or functions, do our epistemic terms serve? We shouldn't be satisfied with as simple an answer as that we use our words like 'knows' and 'rational' in order to refer to knowledge and rational belief. There is a plenitude of things and properties that we could have adopted a word or concept to make specific reference to, but as limited creatures we adopt words and concepts for only some of them. So, something more must be said if we're to explain why we use certain words and concepts rather than others. The explanation might be fairly simple and apparent for words like 'water', 'venomous', 'tiger', or 'Napoleon': clearly it benefits us to be able to think and communicate about such things. But it is

³ See Malcolm (1942), Austin (1946/1961), and Strawson (1952, ch. 9). For useful discussions, see, respectively, Soames (2003, ch. 7), Kaplan (2008), and Avnur (ms.).

⁴ The objection: who cares if you can, or if we did, make up some word that congratulates people who form certain beliefs on the basis of certain evidence? [See Salmon (1957, pp. 39–42).]. The response: if our analysis is correct, it's not just an analysis of some or other congratulatory term. It's an analysis of when a belief is *likely true*, and we care about that. For, being justified/rational/reasonable *just is* being likely true. (This is what's called epistemic or evidential probability. It's not objective chance, the kind found in quantum mechanics—though the two are linked via the Principal Principle.)

⁵ See, e.g. Hookway (2016). The pragmatist tradition has many threads, but I only follow one narrow aspect of it all, namely the idea that Hookway reports like this: 'All pragmatists... have held that the content of a thought or judgment is a matter of the role it fills in our activities of inquiry. The content of a thought or belief is to be explained by reference to what we do with it or how we interpret it.' (Sec. 4.4).

distinctively puzzling why we talk and think about things like knowledge and rationality. An adequate pragmatics must explain why we use these words and concepts (these puzzles are further explored and addressed in the literature that I'll rely on here.⁶).

I want to endorse each of three different views about the pragmatics of our epistemic terms. The views are mutually compatible because our epistemic terms can, and plausibly do, serve all of these multiple functions. This will eventually give us three distinct arguments for anti-skepticism.

(1) The first view comes from Edward Craig.⁷ Craig's view wasn't widely noticed at the time of its publication, but as people are hearing about it they are quickly appreciating it and recently the view has become fairly popular. It's been developed, elaborated, and given new applications by many authors.⁸ Craig's view is that the function of knowledge attributions is to identify, and publicly flag, sources of trustworthy testimony.

It's a fairly intuitive and plausible idea. Our own individual means of collecting information are limited; our own eyes, ears, and reasoning give us only so much information. So, we want to learn more from the eyes, ears, and reasoning of others. But if we gullibly trust just anyone, we're at risk of getting false beliefs. By using knowledge attributions, Craig proposes, we can judiciously sort the trustworthy truth tellers from the rest. The trustworthy sources earn our attribution of knowledge by displaying to us, over time, easily detectible signs of their reliability.

The knowledgable testimony they offer is usually specific to a topic of expertise, and it is most usefully and commonly attributed in the form of *S*-knows-about-*t*, or *S*-knows-whether-or-not-*p*, rather than (the philosopher's favorite form) *S*-knows-that-*p* (which is often redundant to attributors interested in whether-*p*). For example, if I want information about whether there is a gas station nearby, or just what is nearby, you might tell me to ask the locals, the ones who know whether or not there is. When we want information about medical matters we ask the doctor who knows, when we want to know about math we ask the mathematician who knows, when we want to know what's on the menu tonight we ask the chef or the waiter who knows, and so on. When I need certain information, learning who has been flagged as a knowledgable testifier on the matter brings me one step away from the information itself.

(2) The second view about the function of our epistemic terms is one that I've proposed and for which I suggested the name 'epistemic communism'.⁹ Like Craig's view, epistemic communism says that we use epistemically evaluative language in order to make the testimony we rely on trustworthy, that is, safe for consumption. Whereas Craig proposes that we use knowledge attributions to *identify* trustworthy testifiers, communism adds the proposal that our epistemic

⁶ See in particular Dogramaci (2012) for why our use of epistemic evaluation is puzzling.

⁷ Craig (1990).

⁸ See, e.g. Fricker (2007, chs. 5–6), Henderson (2011), Reynolds (2017), Greco and Hedden (2016), and Hannon (forthcoming).

⁹ See Dogramaci (2012, 2015a, b).

evaluations also have the effect of *influencing* others to use the same basic belief-forming methods. We apply negative epistemic evaluations when a testifier is found to have used, whether by a failure in competence or performance, an untrustworthy belief-forming method, perhaps forming a belief by wishful thinking or rushing to judgment or stubbornly refusing to respond to the piled up evidence. By criticizing beliefs as irrational, we derivatively criticize the believer's method, and thereby discourage such poor performance. Likewise, our positive evaluations can reinforce methods that we are eager to see more use of.

When we all form beliefs by using the same basic methods, the benefit is that a testifier's report is then automatically trustworthy. Your testimony is as trustworthy as my own eyes and ears and my own reasoning. This coordination of basic methods is an efficient way of making testimony trustworthy. If a testifier uses an exotic method, it would be unsafe to trust their testimony until the audience has vetted the method as reliable, but vetting is costly and can also leave the testimony redundant. By coordinating, we can each collect more true beliefs while safely handing off to others some of the work of collecting evidence and reasoning.

Craig and I, together, then, propose how our epistemic language (both knowledge attributions and the entailed, or separately issued, justification attributions) makes testimony safe for consumption, efficiently minimizing the risk of gullibly buying false testimony.

(3) The third view about the function of our epistemic terms, this one again focused, like Craig's, just on 'knows', has been influentially defended by Timothy Williamson. The view is that knowledge attributions play an important role in the explanation of behavior.

Williamson's view is a variant of a traditional functionalist view. The traditional view is that our actions are causally explained by our beliefs and desires: our beliefs and desires cause us to take actions that would fulfill those desires if those beliefs were true, and thus our successful actions are often explained by our having true beliefs. Williamson argues that the best explanation of some behaviors must appeal to the person's knowing this or that, rather than merely believing or even truly believing it. He argues that sometimes behavior that is successful, or desire-fulfilling, in a more reliably robust and systematic way, over an extended time, is better explained by knowledge rather than mere true belief. Hilary Kornblith and Jennifer Nagel elaborate and defend the view as well.¹⁰

In Williamson's memorable example (p. 62), a burglar's persistently searching until he finds the hidden diamond may be best explained by his knowledge that the diamond is in the house. An attribution of mere true belief, it's argued, doesn't make for as good an explanation. Positing mere true belief (botched knowledge, as Williamson, p. 47, calls it) wouldn't explain the burglar's successful behavior nearly so well.

I find each of these three views plausible. They're immediately attractive, and they enjoy supportive arguments and defenses. But I leave most of that arguing and defense to other work.

¹⁰ Williamson (2000, esp. secs. 2.4 and 3.4), Kornblith (2002, esp. sec. 2.6), and Nagel (2013).

I take each of these three views about functions to be *teleological explanations* of why we use the words and concepts that we do (rather than the infinitely many alternatives we don't use). They are teleological because they explain why we adopted, or at least why we persist in, our practice by citing the practice's beneficial effects. Explaining a practice by citing its effects raises well known philosophical concerns, but I leave the general defense of the notion of teleology to others.¹¹ (I myself happily follow Millikan and Neander's analysis of teleological explanation as explanation by selection, allowing both natural selection and cultural or artificial selection.) Our three explanations might be incomplete sketches; perhaps there is a larger story that integrates them into natural and cultural evolution. But even as sketches they are still explanatorily illuminating.

My main aim here is application.¹² Taking these three views about pragmatics as premises, I want to argue for anti-skepticism as a conclusion. The next sub-section proposes how to link those premises to that conclusion.

2.2 Linking pragmatics to semantics

How can we draw a link between the way we use words to fulfill a practical function and what our words mean?

I propose the following link between teleological explanations of our use of words, and the meanings those words have—a pragmatics to semantics link.

[P-S] For those words that we can teleologically explain our use of, the full teleological explanation must constrain the words' *contents* such that, if the words normally transmitted beliefs involving those contents, that would explain how our actions serve the words' posited function.

The idea is that pragmatics, as understood here, has to constrain semantics. If our use of one or more words fulfills a function, then the means of doing so must appeal to the meanings, the semantic contents, of those words. And the way semantic content figures in the explanation is via the way beliefs explain actions.

I already brought up, earlier, some important views about how beliefs explain actions. The traditional functionalist says beliefs cause actions that would fulfill the believer's desires if those beliefs were true. We saw that the functionalist and Williamson disagree over whether mere true belief always best explains desire-fulfilling action, but [P-S] is neutral over this dispute. [P-S] simply appeals to the plausible claim that all sides agree to, namely that different beliefs cause different actions, and what makes that difference in belief is a difference in content.

What can be said in support of [P-S]? I hope the general thesis looks plausible on its face. Let me take a paragraph to unpack the thesis a bit in a way that may further

¹¹ A classic defense is Wright (1976). See Sober (2000, sec. 3.7) for a general expository discussion of the use of natural and artificial selection to give teleological explanations. See Millikan (1984), and Neander (1991) for highly plausible analyses of the notion of 'having a function' along the lines of 'having been selected for'.

¹² A similar attempt to merely apply these ideas is the application to the uniqueness debate by Greco and Hedden (2016) and Dogramaci and Horowitz (2016).

bring out that plausibility. [P-S] concerns pragmatics, which consists in teleological explanations of our use of language. Since these explanations are teleological, they must appeal to the *effects* of language use. Those effects are *actions* caused by our use of language, actions that serve a function. This observation raises a question: can such a teleological explanation of our use of words be given without constraining the meanings of those words at all? I don't see how. I don't see how *language* use can cause those function-serving actions unless it's partly by virtue of the used language having particular contents. Of course, our use of sounds (or written marks) can, sometimes, cause actions in a way that's completely independent of any contents they might carry (e.g. I might get your attention by yelling), but then those sounds' status as words, as language, is not any part of the explanation of our behavior here. [P-S] concerns explanations of why we use certain words, why we use language. In a teleological explanation of why we use certain words, those uses are shown to serve a function, but we explain why we use *those words* (as opposed to just explaining why we make certain sounds) only if these uses are taken to be meaningful uses of language, that is, only if they cause us to act one way or another partly because our sounds are carrying distinctive contents. What [P-S] says is that if we employ certain words because they serve a certain function, then we serve that function by using those words *meaningfully*, using them as language and not as mere sounds. Correctly understood, then, I think that [P-S] should be denied only by those who are radical skeptics about content and its explanatory power.

Alongside those general reflections that I hope bring out its plausibility, [P-S] also gets strong support from case studies of meaningful communication. I think the clearest and most compelling examples are the signaling games that David Lewis famously described and studied in his work on convention.¹³ What I borrow from Lewis is his compelling portrait of how, when sounds and marks are used to generate a regularity in behavior that solves a game-theoretic coordination problem, then those sounds and marks intuitively have certain meanings. Those sounds and marks become signals, as Lewis says.¹⁴

Some of the most compelling examples of these signaling games are ones that Brian Skyrms has explored.¹⁵ These are signaling games that naturally evolved in social animals like monkeys, birds, and bees, and even among bacteria and algae. Like Lewis, Skyrms brings out how certain behaviors, initially described in thoroughly naturalistic terms, are intuitively meaningful, and furthermore possess specific meanings that we can intuitively specify. In the famous signaling system of

¹³ See Lewis (1969, ch. 4).

¹⁴ As it's put by Lewis (1969, pp. 124–125): "I have now described the character of a case of signaling without mentioning the meaning of the signals: that two lanterns meant that the redcoats were coming by sea, or whatever. But nothing important seems to have been left unsaid, so what has been said must somehow imply that the signals have their meaning."

Caution! Lewis proposed to naturalize a representational activity, signaling, by reducing it to certain non-representational facts about conventions, but he didn't propose to naturalize language or representational states more generally in this way. I likewise do *not* propose to naturalize language or mental representation by reducing it to conventions.

¹⁵ See Skyrms (2010).

honeybees, a scout finds a food source, returns to the hive, and performs the waggle dance that signals to the hive the direction and distance of the food source.¹⁶ In several monkey species, individuals use acoustically distinctive sounds as alarm calls for distinct purposes. Most famously, this is the practice of the vervets studied by Dorothy Cheney and Richard Seyfarth in the 1980s.¹⁷ When a vervet spots an eagle, a leopard, or a python, it emits a distinct alarm. The eagle alarm causes hearers to hide in the grass, the leopard alarm causes them to run up a tree, and the python alarm causes the monkeys to stand tall and scan the surrounding grounds. These responses are the optimal defenses for each type of predator.

It's apparent to the ethologists studying them, and to anyone who hears the facts, that the different signals used by bees and vervets carry different meanings. The distance the scout bee waggles represents the distance to the food source; orientation relative to the apex of the beehive represents direction relative to the sun. The three vervet alarms represent the three predators. The ethologists studying these animals simply make the intuitive attributions, and readers all find the attributions plausible. What guides these intuitive attributions? We're not conversing with these animals, of course. What enables us to make these attributions is the idea behind [P-S]: if certain signals are made because they serve certain functions, that constrains their meanings. Cheney and Seyfarth explicitly tell us this is their reasoning. "We have called the vocalizations of vervet monkeys semantic signals and have drawn an analogy to human words because of the way these calls *function* in the monkeys' daily lives."¹⁸ While Lewis and Skyrms develop detailed theories about how such signals get their particular meanings, [P-S] only makes the fairly modest claim that the signals carry meanings, and cause beliefs in hearers, that, in virtue of their contents, would explain how the participants serve the practical function that explains the persistence of the signaling system.

I'm next going to argue for anti-skepticism using [P-S], but as we proceed, we'll need to be cautious about the following: [P-S] doesn't require that a word normally communicates *true* beliefs in order to fulfill its function. We need to consider, in each case, whether the best teleological explanation of our practice has us conveying truths. We can defeasibly expect this to be case for most ordinary words, though. Consider a word like 'water'. Its pragmatics are pretty clear, at least in very rough outline. We use this word in order to be able to think and talk about a biologically very important substance, water. But we don't just want any old sentences or thoughts, we want true ones of course (and not only true, but also relevant to our current practical predicament). So, the full pragmatics, the full teleological explanation of why we use 'water', will have it that, in fulfilling the word's function, we do normally use it to communicate true beliefs. But it had better not be that every word we can teleologically explain our use of is normally used to convey truths. Our lives are filled with lies and myths. I don't know exactly what

¹⁶ See von Frisch (1967); summarized by Skyrms (2010, p. 28).

¹⁷ See Cheney and Seyfarth (1990). Skyrms (2010, p. 22) lists several other monkey species.

You can listen to the vervet alarm calls for yourself here: <http://web.sas.upenn.edu/seymarth/vocalizations/vervet-monkey-vocalizations/>.

¹⁸ Cheney and Seyfarth (1990, p. 139), emphasis in original; see also chs. 1, 4–7.

function our empty language for the supernatural serves, but it certainly seems possible that it serves some function, maybe a biological evolutionary function, even if we don't use the language to systematically speak truthfully. Some examples of false signaling with a clear function are provided by animals. Beavers, some birds, and some monkeys use alarm signals, but emit a false alarm more frequently than a true alarm (unlike most vervets).

2.3 Our epistemic terms have anti-skeptical extensions

Suppose [P-S], and suppose any one of the three views of the pragmatics of our epistemic terms that were sketched earlier. Then, I claim, anti-skepticism follows. Let me elaborate how so.

Take, to begin with, Williamson's explanatory function of knowledge attributions. Suppose it is true: we do use knowledge attributions because they allow us to give explanations. We use 'knows' because it allows us to explain why the burglar persistently searched until he found the diamond. A teleological explanation of some word must connect the word's use to some (somehow beneficial) effected actions. And according to [P-S], that effected action is caused partly in virtue of the *contents* of our explanations by knowledge attribution. What are these effected actions? The ability to explain the past confers a more immediately useful ability to predict the future, and so our effected actions might include things like predicting the burglar will have the diamond on him when he exits the house, and then arresting him red-handed. These sorts of activities can only plausibly persist if our explanations by knowledge attribution are on the mark reasonably frequently. If our knowledge attributing explanations are all, or even just mostly, erroneous, then we've lost our teleological explanation; it fails. We can't reasonably take the effects of *erroneous* knowledge attributions to allow us to teleologically explain why we use 'knows'. So, if [P-S] and the explanatory and predictive function of knowledge attributions are right, then our uses of 'knows' are often enough correct. That is, 'knows', if it is to serve its explanatory function, really has in its extension the burglar's belief about the diamond. That's anti-skepticism semantically formulated, and, by semantic descent, anti-skepticism follows.

I cautioned above that [P-S] doesn't require that our words always convey true beliefs, and [P-S] should allow us to explain our occasional use of empty, non-referring words too. What I'm arguing here, though, is that 'knows' is not like that. Knowledge attributions are not systematically false or empty, because if they were they could not serve their explanatory-predictive function.

In order to fulfill the Craigian function too, knowledge attributions must, frequently enough, be true. On Craig's view, knowledge attributions flag the testifiers whose testimony is then accepted and acted upon. Now, it's fairly clear that, if Craig's teleological explanation of our use of 'knows' is correct, then the *testimony* of the flagged testifiers should normally be true, i.e. it should be reliable: if that testimony resulted in hearers adopting, and acting upon, false beliefs, then the practice would not confer any apparent benefit. But, what anti-skepticism requires is that the *knowledge attributions* be true. Must these be true if Craig's teleological explanation is correct? [P-S] doesn't directly require that our uses of 'knows'

convey truths, however, it does require that the complete teleological explanation assigns 'knows' some content such that the beliefs that our uses of 'knows' produce must *explain* our function-fulfilling actions. Those actions are the successful identification of, and subsequent reliance upon, trustworthy testifiers. And it's hard to see how we could best explain the success, the function-fulfilling power, of these actions unless we are *correctly* identifying knowers. What's the skepticism-permitting alternative? The alternative is that we are correctly identifying trustworthy testifiers and beneficially relying on their reliable testimony (Craig's view requires that much), but it is *false* that these trustworthy testifiers are knowers. The complete teleological explanation here should, like any explanation, be the one that, all else equal, gives the *simplest* accounting of all the data, the data in this case being how we use 'knows' so as to fulfill Craig's function. The preference for simplicity applies to the semantic content our explanation assigns to 'knows'. And here it seems it would only be gratuitously complex to have the explanation give 'knows' a content that makes our knowledge attributions false, even while we handily make use of those attributions to identify and place our trust in reliable testifiers. So, on the basis of Craig's proposal, [P-S], and the virtue of giving the simplest explanation (all else equal), I take this to amount to a second argument for anti-skepticism.

Our first two arguments for anti-skepticism support the anti-skeptical conclusion, but without directly addressing the skeptic's own argument for skepticism. Earlier I said the key premise in the skeptical argument, or family of arguments, is the evidence neutrality premise (which says, recall, that all the evidence we can ever get doesn't favor the ordinary propositions we believe over skeptical hypotheses, like that we're brains in vats). We can now give not only a third argument for anti-skepticism but also an argument against evidence neutrality, by turning to the remaining function I've posited for epistemic language, my communist function.

Epistemic communism says our epistemic evaluations have an influence on our audiences. They promote the use of trustworthy belief-forming *methods*. Communism recognizes that epistemic evaluations don't target lone beliefs, they target belief-forming methods, which relate beliefs to evidence.¹⁹ A belief is justified iff: the believer's evidence supports it and she followed a method that had her base that belief on that evidence. So, our positive and negative evaluations of beliefs (as justified or unjustified) are sensitive to the relation we express with the word 'evidence', namely the *x-is-evidence-for-y* relation. Communism says the function of epistemic evaluations is to promote the use of certain methods, and thus also, in effect, to promote the adoption of certain beliefs on certain evidence.

What methods, and in particular what candidate *x-is-evidence-for-y* relation, is it beneficial to promote? Communism says that we promote a shared set of methods so that we don't need to expend resources vetting exotic methods, checking their track-records before we can determine if they're reliable. But it's not beneficial to share just any old methods. The methods also need to be reliable in the actual world. So, communism says that the beneficial effect of our evaluative practice is the

¹⁹ See Dogramaci (2012, sec. 2.1) for arguments.

promotion of methods that are shared and actually reliable. This imposes a constraint on the *x-is-evidence-for-y* relation that our evaluations recommend: *x* needs to be a reliable indicator of *y*, in the actual world. This is what rules out our use of a predicate, ‘evidence’, with a skeptical extension. Since the actual world is not a skeptical world, since for example we’re not brains in vats, we use our epistemic language to express some candidate evidence relation that is anti-skeptical. This is what serves us best, and what a teleological explanation of our use of the language should attribute to us.

This gives us an argument for the contrary of evidence neutrality, and thereby a third argument for anti-skepticism. If epistemic communism is right, we use the language of epistemic evaluation, not only ‘knows’, but ‘justified’, ‘evidence’ and others, for the purpose of promoting certain belief forming methods. These are the methods where we call the basis “good evidence” for the belief that rests on it. The methods, and the accompanying evidence-for relation, that it benefits us to promote are the ones that produce trustworthy beliefs. Trustworthy beliefs must be, among other things, reliable in the actual world. To believe, on our actual evidence, that we’re brains in vats, would be to use an unreliable method. What it benefits us to promote, and what we in fact do promote, are methods that prescribe our ordinary beliefs on the basis of our ordinary evidence. By [P-S] and the same preference for simplicity we used above, our epistemic terms which serve this function, including ‘knows’, ‘justified’, and ‘evidence’, have anti-skeptical extensions. And thus, by a step of semantic descent, our evidence favors, and we can know, that we’re not brains in vats.

This last argument concluded that ‘evidence’ expresses some candidate anti-skeptical extension. Earlier I said we have several examples of what such a candidate extension could be, examples in the forms of the extensions that would make true an anti-skeptical theory such as dogmatism, neo-rationalism, or anti-closure. To be clear, the argument I’ve given is neutral among such options. It is an argument for the disjunction of particular anti-skeptical theories. Nevertheless, it’s an argument for anti-skepticism.

The arguments I’ve just offered have many premises. Not only do they have premises about the function of epistemic language, and [P-S], the arguments make assumptions about what kinds of predictions and methods are reliable in the actual world, and consequently they assume, among other things, that we’re not brains in vats. These assumptions require empirical justifications, but we’re using them to argue against, for one thing, skepticism about empirical justification. Is an argument for anti-skepticism that rests on such assumptions any good? The next section considers what good can come from such an argument.²⁰

²⁰ Craig (1990) and Reynolds (2017) don’t propose to combat skepticism (though Craig (chs. 12–13) explores where skepticism’s appeal might originate from). Others have considered using the Craigian view to argue against skepticism, though none resembles my approach here; see Hannon (forthcoming), Fricker (2008, sec. 3.3), and a brief remark in Gardiner (2015, p. 42).

3 The varieties of anti-skeptical arguments

In this section, I distinguish different kinds of anti-skeptical arguments we might make, and I place the arguments of this paper in a novel category. I then briefly explore the benefits of this kind of argument.

3.1 Anti-skeptical projects: modest, impossibly ambitious, reasonably ambitious

Pryor (2000, p. 517) drew a two-part distinction between modest and ambitious anti-skeptical projects. Both aim to secure the claim that we have plenty of ordinary knowledge. The difference concerns what assumptions they grant themselves, and thus also what anti-skeptical work they see as needing to be done.

The modest project helps itself to a wide range of assumptions. In executing the project, any arguments you make can take for their premises anything you know at the outset of inquiry. This includes your knowledge of anti-skepticism, i.e. your knowledge that you have plenty of knowledge. That is, the modest project grants itself anti-skepticism as a reasonable assumption. Consequently, the modest project is one that's not concerned with making particularly interesting arguments with anti-skepticism as their conclusion. The work done by the modest project is rather defensive work. Since the skeptic argues against something we believe, we need to deny one of the skeptic's own premises, revising own overall views as necessary. So, as Pryor observes, the modest project still requires some very hard work, since it's hard work to diagnose where the skeptic's own arguments and their tempting premises go wrong. Pryor engages in the modest project when he proposes dogmatism as the correct anti-skeptical view, as the view that best preserves and revises our initial views while diagnosing the skeptic's errors.²¹

The ambitious project is to argue *for* anti-skepticism—i.e. argue that we have a fair amount of knowledge—while sticking to premises that aren't disputed by the skeptic. Exactly which premises are off limits? We can answer that question in different ways (Pryor's own characterization leaves open either interpretation of the ambitious anti-skeptical project (see pp. 517–8)).

On a less interesting interpretation, what's off limits is *every premise that the skeptic argues we don't know*. So, we can't argue from, e.g. the premise that we have hands. Meeting these standards seems most ambitious. For one thing, it would be logically impossible to meet these standards in arguing against a wholesale skepticism that denies us all knowledge.²² It's not my project.

²¹ See Pryor (2000, pp. 518, 547, and fn. 37).

²² See Rinard (forthcoming) for an argument that is explicitly as ambitious as this, though targeting only external world skepticism. See also Vogel (1990) for an argument that can be seen as meeting these ambitious standards; Vogel also targets only external world skepticism, a traditional Cartesian sort of external world skepticism that grants us introspective knowledge. Rinard argues that external world skepticism is self-undermining, and Vogel aims to use IBE to rest our external world knowledge on our introspective knowledge. Unlike them, I help myself to premises concerning the external world.

On what I think is a more interesting interpretation, what's off limits are just the premises the skeptic *denies*. The skeptic does *not* deny any *non-epistemic* claims, e.g. he does not deny that we have hands. What the skeptic denies is just knowledge and justification attributions. So, on this interpretation, it's just knowledge and justification attributions that are off limits (as well as any equivalent epistemic claims). This gives us something that's in between the impossibly ambitious anti-skeptical project and the modest project that Pryor engaged in. This project is reasonably ambitious.

The arguments for anti-skepticism I've given here are in this middle range of ambitiousness. My project is not the most ambitious one, because I do help myself to non-epistemic assumptions about things like what explains certain behaviors and what sort of testimony is reliable. But my project is more ambitious than the modest project of Pryor, because my arguments avoid resting on any epistemic premises, in particular any knowledge attributions. My arguments treat these as off limits.

I take this to make them arguments that, in an interesting sense, don't beg the question; they are non-trivial, non-circular arguments for anti-skepticism. The arguments *conclude* that we have plenty of knowledge, but *no premise* attributes us any knowledge. So, in an intuitive sense the arguments conclude with anti-skepticism but they don't *assume* anti-skepticism.

My arguments will look as if they assume anti-skepticism to someone who thinks an argument assumes not only its premises but also the further claim *that you know* its premises. I don't find this view of what an argument "assumes" appealing. It looks to me arbitrary and unmotivated. I suspect its appeal (to others) is due to the appeal of the highly controversial KK (you know what you know) principle on which having any single piece of knowledge requires you to have infinitely many pieces of higher-order knowledge.

3.2 The possible value of "reasonably ambitious" arguments for anti-skepticism

My arguments do assume we're not brains in vats. I assume we are in a real world full of successful actors and reliable testifiers. I assume that our actual ordinary belief-forming methods are by-and-large reliable. But this does not make my arguments question-begging, or valueless, as I'll try to show now.

I can see four possible ways in which a reasonably ambitious anti-skeptical argument can have distinctive value.

(1) Perhaps someone bamboozled into skepticism would have a use for my arguments, using them to reason their way to anti-skepticism from non-epistemic premises that the skeptic did not deny. This is an advantage my arguments have, and it's an advantage they have over the modest project which offers no help to skeptics themselves.²³ For some skeptics, this might be a limited benefit my arguments offer: even though no skeptic denies my premises, the evidence neutrality principle may leave a skeptic with a middling credence in these premises, and, if so, then their

²³ See Pryor (2004, sec. 7, and the cited authors in his fn. 47) for concessions of this point.

rational credence in the conclusion *might* only get dragged up to a similarly middling degree. (I italicize ‘might’ here because the issue of when a boost in the credence in some premises can boost credence in a consequent conclusion is complex and controversial.²⁴) My arguments could be more helpful, however, to another sort of victim bamboozled into skepticism, a victim who did deny knowledge attributions but managed to retain their justified beliefs and confidence and perhaps even their knowledge of non-epistemic facts. Such a victim knows they have hands and that tables and chairs exist, but doesn’t know they have any knowledge. By reasoning from the premises of my arguments to the conclusion of anti-skepticism, such a person could come to know anti-skepticism. This reasoner could even be following the very demanding knowledge-norm that requires reasoners to only believe and reason from premises they know. I could not hope to convince any skeptic *that they will be obeying the norm* if they should engage in such reasoning, but my suggestion here is rather that certain skeptics *could reason* in this way and thus come to inferentially know anti-skepticism from non-epistemic premises that were known.

This sort of reasoning that our bamboozled skeptic could follow would be what DeRose (2000, pp. 129–130) calls ‘heroic’: an anti-skeptical strategy is heroic if it aims to *gain* us the knowledge that the skeptic argues we lacked. DeRose suggests that epistemologists shouldn’t bother trying for heroism. I concede that my suggested path to heroism at best could help only very few actual people. I certainly don’t think *I* need to reason my way out of skepticism: I am confident that I’ve always had plenty of knowledge, and I’m more confident of that than I am of the premises of any argument for anti-skepticism. But I have had two famous epistemologists tell me in person that they are genuine skeptics—they believe there are no instances of knowledge. And of course a perplexing number of undergrads also claim to be skeptics. So, I count this as one possible value, however significant, for the reasonably ambitious anti-skeptical project.

(2) Furthermore, a reasonably ambitious argument for anti-skepticism could still have some added value even for those of us who already rationally believed, or even knew, that anti-skepticism was true before doing any philosophy. The argument could shore up our justified confidence in anti-skepticism. Our confidence could rationally increase, or gain new resiliency against some potential defeaters. This is a small but reasonable dose of heroism.

(3) Another value of a reasonably ambitious argument might be that it helps clarify the range of possibilities anti-skepticism holds true in. It might draw new distinctions within that range. Here’s a possible illustration. Above I highlighted that the anti-skeptic must assign an anti-skeptical extension for the *x-is-evidence-for-y* relation. One way or another, this extension needs to include evidence we actually possess paired together with the ordinary propositions we know. We might notice now that anti-skeptics have some flexibility over how this can go. We can be generous with how much to include in either the first or rather the second relatum of the relation. We can either say that we have lots more evidence than skeptics might

²⁴ See Kotzen (2012) for discussion. Kotzen’s own view would support the worry I say might arise here.

allege, or we can say that a slimmer stock of evidence can still justify all the ordinary propositions we know. (In rough caricature, Williamson's view might illustrate the first option, and Descartes's might illustrate the second.) Ambitious anti-skeptical arguments might lend us guidance as we tinker over what to say is the extension of the *x-is-evidence-for-y* relation.

I won't defend this here, but my inclination is to see the communist and Craigian arguments as encouraging us to just pack more into the first relatum. That is, the role those views give to our epistemic concepts suggests to me that it is easier to acquire tons of good evidence than skeptics would have us think, rather than suggesting that evidence is more powerful, and less neutral, than skeptics would have us think. But, again, I won't argue for these ideas here, since I'm just sketching how ambitious arguments can possibly have value.²⁵

(4) Finally, the value of an ambitious argument that we might be most hopeful for is its potential to provide new and deeper explanatory insights. An ambitious argument might, to some extent or in some new way, help explain why anti-skepticism is true. I hope it's somewhat intuitive that reflections on ordinary language and its functions can be explanatorily illuminating, including of why we have knowledge. I do hope that's intuitive—I'm not presupposing any theory of explanation. I hope the arguments for anti-skepticism themselves, if they are taken to be good arguments, demonstrate the possibility.

I'll admit I have a bit of inclination to accept a general view on which *any* new way to deduce from known premises a conclusion that you already know helps you to better explain that conclusion, in some sense and to some extent. Some deductions are famously poor explanations, but even they might nevertheless still add some little bit of explanatory value.

I think it's intuitive that even an argument that takes a step of semantic descent, as the arguments in this paper do, can yield an explanation, though such arguments don't meet most textbook definitions of 'deduction'. A simple example of such an explanatory argument: Why are no lies known? Because 'knows' is factive. I think we can, to some degree, intuitively explain why no lies are known by "deducing" that from that semantic premise about 'knows'. Some philosophers may think we cannot explain such a non-semantic fact by such a semantic explanation; that (incorrect) thought may be based on the (correct) point that semantic facts do not generally *make* non-semantic facts true. The semantics of 'bachelor' does not *make it the case* that bachelors are unmarried.²⁶ But despite that, I think there remains a highly intuitive semantically formulated explanation, one we actually frequently give in ordinary life, of why bachelors are unmarried: it's *because* that's part of what 'bachelor' means. Perhaps the best example in support of this possibility (of semantic, ultimately pragmatic, facts explaining non-semantic facts) is Craig's own work. In his original book, Craig (1990) used his thesis about the function of 'knows' to predict and explain many features that epistemologists found knowledge

²⁵ See Miller (2016) for why this is viable. He shows that fans of both dogmatism and Bayesianism can say we perceptually acquire propositions like ⟨I have a hand.⟩ as our evidence.

²⁶ See Boghossian (1996).

itself to have or lack: the claim that ‘knows’ flags acceptable testifiers explains why knowledge must be factive and reliable, and it also explains the falsity of the theory that reduces knowledge to reliably formed true belief (and Craig explains the falsity of several other popular reductive theories too).

An ambitious argument might also lend explanatory insight by helping to show where certain arguments for skepticism go wrong. An example of this would be my earlier suggestion that communism is better equipped than the other views to explain why the evidence neutrality premise is wrong. This value is one that both the modest project and my ambitious project make a claim to. On Pryor’s development of Moore’s idea for explaining where the skeptical argument goes wrong, for example, the skeptic goes wrong when he says we can’t know we’re not brains in vats. Pryor and Moore say we *can* know that *because* we know we have hands and this is a sound basis for knowing we’re not brains in vats. Ambitious arguments can offer additional ways to explain why the skeptic’s argument goes wrong, again for example the communist explanation of why the evidence neutrality premise is false.

But even aside from that defensive task of explaining where the skeptic’s argument goes wrong, I want to emphasize the distinctive explanatory value I see in the positive work of arguing, reasonably ambitiously, for anti-skepticism. Why is it true that we have lots of knowledge? Here is something we can say that intuitively helps explain why: it’s true because of the way we use our epistemic words and concepts. Strawson said, concerning our anti-skeptical usage of our epistemic language: “Doing this is what ‘being reasonable’ means in such a context” [Strawson (1952, p. 257)]. As I want to put it, though: making true knowledge attributions is what we use these words *for*.

Acknowledgements For their help with this paper I’d like to thank audiences at Arizona State University, Georgia State University, the University of Miami, the University of Nebraska, TEX (the Texas Epistemology Xtravaganza), and the 2018 Pacific APA where Jennifer Nagel gave valuable comments. I’m especially grateful to Yuval Avnur for extensive discussions about the ordinary language argument.

References

- Austin, J. L. (1946/1961). Other minds. In J. O. Urmson & G. J. Warnock (Eds.), *Philosophical papers* (pp. 76–116). Oxford: Oxford University Press.
- Avnur, Y. (ms.). Justification as a loaded notion.
- Boghossian, P. (1996). Analyticity reconsidered. *Nous*, 30(3), 369–392.
- Cheney, D., & Seyfarth, R. (1990). *How monkeys see the world: Inside the mind of another species*. Chicago: Chicago University Press.
- Craig, E. (1990). *Knowledge and the state of nature*. Oxford: Oxford University Press.
- DeRose, K. (2000). How can we know that we’re not brains in vats? *The Southern Journal of Philosophy*, 38(Supplement), 121–148.
- Dogramaci, S. (2012). Reverse engineering epistemic evaluations. *Philosophy and Phenomenological Research*, 84(3), 513–530.
- Dogramaci, S. (2015a). Communist conventions for deductive reasoning. *Nous*, 49(4), 776–799.
- Dogramaci, S. (2015b). Forget and forgive: A practical approach to forgotten evidence. *Ergo*, 2(26), 645–677.
- Dogramaci, S., & Horowitz, S. (2016). An argument for uniqueness about evidential support. *Philosophical Issues*, 26, 130–147.

- Dretske, F. (1970). Epistemic operators. *The Journal of Philosophy*, 67(24), 1007–1023.
- Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford: Oxford University Press.
- Fricker, M. (2008). Scepticism and the genealogy of knowledge: Situating epistemology in time. *Philosophical Papers*, 37(1), 27–50.
- Gardiner, G. (2015). Teleologies and the methodology of epistemology. In D. Henderson & J. Greco (Eds.), *Epistemic evaluation: Purposeful epistemology* (pp. 31–45). Oxford: Oxford University Press.
- Greco, D., & Hedden, B. (2016). Uniqueness and metaepistemology. *The Journal of Philosophy*, 113(8), 365–395.
- Hannon, M. (forthcoming). *What's the point of knowledge?* Oxford: Oxford University Press.
- Henderson, D. (2011). Gate-keeping contextualism. *Episteme*, 8(1), 83–98.
- Hookway, C. (2016). Pragmatism. In E. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <https://plato.stanford.edu/archives/sum2016/entries/pragmatism/>.
- Kaplan, M. (2008). Austin's way with skepticism. In J. Greco (Ed.), *The Oxford companion to skepticism* (pp. 348–371). Oxford: Oxford University Press.
- Kornblith, H. (2002). *Knowledge and its place in nature*. Oxford: Oxford University Press.
- Kotzen, M. (2012). Dragging and confirming. *The Philosophical Review*, 121(1), 55–93.
- Lewis, D. (1969). *Convention: A philosophical study*. Cambridge, MA: Harvard University Press.
- Malcolm, N. (1942). Moore and ordinary language. In P. A. Schilpp (Ed.), *The philosophy of G. E. Moore* (Vol. 1, pp. 343–368). La Salle, IL: Open Court.
- Miller, B. (2016). How to be a Bayesian dogmatist. *Australasian Journal of Philosophy*, 94(4), 766–780.
- Millikan, R. G. (1984). *Language, thought and other biological categories*. Cambridge, MA: MIT Press.
- Nagel, J. (2013). Knowledge as a mental state. *Oxford Studies in Epistemology*, 4, 273–308.
- Neander, K. (1991). The teleological notion of 'function'. *Australasian Journal of Philosophy*, 69(4), 454–468.
- Pryor, J. (2000). The skeptic and the dogmatist. *Nous*, 34(4), 517–549.
- Pryor, J. (2004). What's wrong with Moore's argument? *Philosophical Issues*, 14, 349–378.
- Reynolds, S. (2017). *Knowledge as acceptable testimony*. Cambridge: Cambridge University Press.
- Rinard, S. (forthcoming). Reasoning one's way out of skepticism. In K. McCain & T. Poston (Eds.), *The mystery of skepticism*. Leiden: Brill.
- Salmon, W. (1957). Should we attempt to justify induction? *Philosophical Studies*, 8(3), 33–48.
- Skyrms, B. (2010). *Signals: Evolution, learning, and information*. Oxford: Oxford University Press.
- Soames, S. (2003). *Philosophical analysis in the twentieth century, volume 2: The age of meaning*. Princeton, NJ: Princeton University Press.
- Sober, E. (2000). *The philosophy of biology* (2nd ed.). Boulder: Westview.
- Strawson, P. F. (1952). *Introduction to logical theory*. London: Methuen.
- Stroud, B. (1984). *The significance of philosophical scepticism*. Oxford: Oxford University Press.
- Vogel, J. (1990). Cartesian skepticism and inference to the best explanation. *The Journal of Philosophy*, 87(11), 658–666.
- von Frisch, K. (1967). *The dance language and orientation of the bees*. Cambridge, MA: Harvard University Press.
- White, R. (2006). Problems for dogmatism. *Philosophical Studies*, 131, 525–557.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Wright, L. (1976). *Teleological explanations*. Berkeley: University of California Press.
- Wright, C. (2004). Warrant for nothing (and foundations for free)? *Aristotelian Society Supplementary*, 78(1), 167–212.