

TELEOLOGICAL EPISTEMOLOGY*

Jane Friedman | jane.friedman@nyu.edu

01/2018

1 Introduction

It is typically thought that some epistemic states are valuable – knowing, truly or accurately believing, understanding (to name a few). These are states it's thought good to be in and it's also said that we aim or want to be in them. It is then sometimes claimed that these sorts of thoughts about epistemic goods or values ground or explain our epistemic norms. For instance, we think subjects should follow their evidence when they form their beliefs. But why should they? Why not believe against the evidence or ignore it completely in deciding what to believe? Here's a compelling sort of answer: because epistemic subjects are or ought to be trying to have true beliefs or to know and following their evidence is a means to ends like those or to fulfilling obligations like those.

Let's call norms that are grounded in or explained by values – either valuable states or states of affairs or the things that actual epistemic subjects value – 'teleological'. I am going to focus the discussion to come on a natural way of understanding the idea that our familiar epistemic norms are teleological and argue that on this way of understanding the idea, they are not teleological. And as we'll see I think the sort of argument I make can be extended to other sorts of teleological accounts of our epistemic norms as well.

2 Clarifying the issues

There is a significant teleological thread running through contemporary epistemology.¹ The idea that some doxastic or epistemic states are valuable is entirely commonplace. As is the thought that epistemic rationality is somehow tied to the pursuit of these goods.² There are a number of different epistemic states

*Please cite published version: Friedman, J. (2019), TELEOLOGICAL EPISTEMOLOGY. *Philosophical Studies* 176: 673–691. <https://doi.org/10.1007/s11098-018-1033-7>.

¹For a nice accounting of this thread in its many different forms see Talbot (2014), p. 602.

²One sort of place in which this thought comes out: virtually every theory of justification aims to have it be the case that justified beliefs are more likely (in some sense of 'likely') to be true than, say, lucky guesses.

that have been thought of as valuable and a number of different ways of fleshing out the thought that rational epistemic subjects are subjects trying to get into these valuable states.

In order to have a clear target here, I'm going to focus on one way of precisifying this teleological thread in epistemology. I don't know that anyone has explicitly endorsed my exact precisification, but I think that the sorts of arguments I am going to make here can be made for others as well. In what's to come I'm going to focus on a teleological epistemology that makes knowledge the primary epistemic good or end. I'll call this 'teleological_k epistemology' and its norms 'teleological_k norms'. When epistemologists have discussed our epistemic goals or values the focal goal or value has often been truth or true belief rather than knowledge.³ Let me say at the outset then that the arguments I make here could just as easily be made against a teleological_t epistemology (one according to which the relevant valuable/valued end is true belief).⁴ Personally, I find the thought that our aim is to merely believe the truth rather than know it implausible, which is why I'm staying focused on teleological_k epistemology.⁵

Moreover, I am going to understand the claim that rational subjects are subjects in pursuit of epistemic ends or values as the claim that the norms of epistemic rationality are the norms that subjects in pursuit of the relevant ends – in this case knowledge – ought to conform to. My thought here is that the claim that epistemic rationality is somehow a matter of goal pursuit can be thought of as the claim that our familiar epistemic norms are the norms that rational subjects trying achieve the relevant goal (e.g., to know more) are going to follow. Near the end of the paper I'll say a bit more about other ways of thinking about teleological epistemology.

All of this means that we can see the teleologist_k as making two main claims:

- (1) A value claim: knowledge is valuable/valued.

³For some explanation for this focus see David (2001).

⁴More recently, a teleological account of the norms of probabilistic coherence has become popular as well (largely inspired by Joyce (1998)). In this case the relevant goal or value is credal accuracy. I don't think that the arguments I make here bear directly on these sorts of views. I'm focused here on norms for belief rather than the sorts of norms for credences that are the focus of those discussions.

⁵An important articulation of the teleological structure of normative epistemology comes to us from William James who famously claimed that we had twin goals as rational epistemic subjects: getting to truth and avoiding error. James (1896) was explicit about the thought that the truth goal is the goal of knowing the truth.

- (2) A grounding claim: that knowledge is valuable/valued grounds or explains our central epistemic norms; these norms are ‘teleologically_k grounded’.

The focus of this paper is largely on (2); I want to argue that (2) is false. Before we get to (2) though let me first make a couple of points about (1).

First, in epistemology we often talk as if states of mind or attitudes are the bearers of value. Sometimes though it’s easier to talk about states of affairs as having or lacking value. I’m not going to make anything of this distinction here, but assume that we can easily translate one sort of talk into the other, i.e., that to say that one’s knowing p is valuable is to say that the state of affairs in which one knows p is valuable, and vice versa. I’ll talk in both ways in what’s to come.

Second, there are two ways of thinking about the claim that knowledge is valuable that we should distinguish – one more “subjective” and one more “objective”. The subjective claim is that knowledge or having knowledge is something that people in fact value – this is a claim about what actual people care about. The objective claim is the claim that knowing is valuable – this is a claim about some state having value, and can be independent of any claims about what actual (or possible) people care about or value. This distinction raises a complex cluster of issues, but we don’t need to engage with them here. I mean the arguments to come to have their impact given either way of thinking about the claim that knowledge is valuable.

Now more on (2). What is it for a norm to be teleologically grounded? I’m thinking of a norm as expressed by a sentence of the form ‘ Δ S φ s’ or ‘ Δ S φ s in c ’, where ‘S’ should be replaced with the name of a subject ‘ Δ ’ with a phrase expressing some normative modality, ‘ φ ’ a phrase picking out some state or action, and ‘ c ’ the specification of some relevant circumstances. ‘ Δ ’ can express any normative modality on the way I’m thinking about norms here, e.g., ‘ Δ ’ can be replaced with phrases like ‘It ought to be the case that’, ‘It’s permitted to be the case that’, ‘There is reason for it to be the case that’ and so on. ‘ φ ’ in the norms we’ll think about will have to do with forming and having beliefs.

So each norm has a central normative modality. We might ask what explains why subjects are normatively bound in the way some norm claims – why subjects ought to do something or be some way (in some circumstances), or why they are permitted to do something or be some way (in some circumstances).⁶ My

⁶On ‘bound’: I want to say that a norm binds a subject so long as it applies to the subject.

thought is that if a norm is teleologically grounded, then that explanation will go by way of some valuable or valued state of affairs. There are several ways this sort of explanation can go, but I want to stay focused on a central way it tends to go: it will be the case that someone (say) ought to (say) do something because doing that thing will or is likely to lead to the relevant valuable or valued state of affairs. The general form of the sort of explanation that I'm interested here then is that subjects are normatively bound to do something or be some way because doing that thing or being that way conduces to the relevant valuable state of affairs. For instance, subjects ought to follow their evidence because doing so is likely to get them more knowledge.

It's worth relating the idea that epistemic norms are teleologically grounded to some others. First, there has been some discussion recently about "epistemic consequentialism".⁷ It's hard to say exactly what makes some view consequentialist, but one thought is that some normative theory is consequentialist if, according to the theory, normative facts of one kind, e.g., deontological facts, facts about what we ought or are permitted to do, depend on or are grounded in facts about the value of states of affairs brought about in action. For instance, on views like this it's exactly because some states of affairs are valuable that we are required (or permitted, etc.) to act in ways that help bring them about. But then take some norm n that says that one ought to φ in c . Plausibly, we can think of n as a consequentialist norm if what explains why one ought to φ in those circumstances is the value of some state of affairs one will bring about by φ -ing in those circumstances. But then thinking of n as a consequentialist norm means thinking of it as teleologically grounded. So this paper can be thought of as containing an argument against at least some forms of epistemic consequentialism.

Relatedly, forms of "epistemic instrumentalism" are targets here, too. There is a debate in contemporary epistemology about whether our epistemic norms are instrumental norms: whether what explains why we should conform to them is some sort of epistemic goal we have (or at least should have). For instance, should we be following our evidence because that's a means of achieving some

Sometimes it's said that only requirement norms bind and others apply in other ways, but I'm not making that distinction here.

⁷For some of this see, Berker (2013a), Berker (2013b) and Ahlstrom-Vij and Dunn (2014). Also see the papers in ?.

epistemic goal or ours, e.g., having more true beliefs or knowledge?⁸

As far as I can tell, for the claim that our familiar epistemic norms are instrumental norms to be at all plausible, it needs to be claimed that our epistemic goals are very general or global, rather than local or tied to specific subject matters.⁹ The injunction to respect the evidence isn't a demand to respect it with respect to some subject matters but not others. Only the general or global aim of believing the truth or knowing (whatever the subject matter) stands a chance of explaining our familiar epistemic norms.

How does the explanation by way of these sorts of global goals go? The thought is that it's because we all have (or at least should have) the general end of believing the truth or of knowing that we should, e.g., follow our evidence. So if epistemic rationality is a form of instrumental rationality then what explains why one should respect or follow one's evidence is that respecting or following one's evidence is a means to achieving one's epistemic goals. In general then it's because doing the relevant thing conduces to some valued end that one ought to do the relevant thing.¹⁰ This sort of epistemic instrumentalism is a target here as well then. Even if we do or should have these sorts of global epistemic goals, the argument here will show that they don't ground our familiar epistemic norms; those norms are not instrumental norms in the relevant sense.

Apart from these two families of teleological positions, there is something intuitive about a teleological epistemology. Shouldn't there be some explanation as to why we ought to follow our evidence? Or why we're permitted to? Or why we should use reliable rather than unreliable methods? Plausibly there should be and the idea that the answer has something to do with following our evidence and using those methods as conducing to a state of value or that we value has great appeal.

In fact, Talbot (2014) argues that the thought that our familiar epistemic norms are grounded in or explained by (Talbot talks of their "arising from") epistemic ends or values is widely accepted across epistemology. Talbot him-

⁸For some of the instrumentalism debate in epistemology see Foley (1987), Stich (1990), Kornblith (1993), Kelly (2003), Leite (2007), Grimm (2008), Street (2009) and Steglich-Petersen (2011). See Cowie (2014) and ? for a couple of recent defences.

⁹ Foley (1987) explicitly endorses a picture like this. According to Foley epistemic rationality is a matter of our properly pursuing our global epistemic goal. Foley describes the goal as, "concerned with now believing those propositions that are true and now not believing those propositions that are false." See p. 8.

¹⁰See Feldman (2000) for a defence of evidentialism in this spirit.

self is focused on the goal of believing the truth rather than knowing it, but as I've said, I don't think that difference makes much of a difference to what's to come. Talbot's thought is that it is extremely common to deny that our epistemic rules are "brute", to claim instead that they arise from some source, and finally to think of that source as an end or goal or value. Talbot argues that this feature of epistemic normativity means that we should accept some controversial epistemic norms. This seems plausible to me. What I want to do though is think about whether this sort of teleological picture can even do the job it's being tasked with: explaining or grounding our familiar (and largely uncontroversial) epistemic norms. I think it cannot.

It is not my plan to argue by way of counterexample or counterintuitive consequences here. Instead my plan is to think about the set of teleologically_k grounded norms and show that many familiar epistemic norms won't be in that set.

3 Teleological trouble

To start let's think of normative epistemology "as we know it" or as it is currently done as a set of norms, E. Of course there is debate about which norms are in E, but nothing to come will hang on putting anything too controversial in E or leaving anything too beloved out. I'll say more about E in a moment, but in general it should contain our familiar epistemic norms. I also want to think about another set of norms, T. I want to think of T as an exhaustive set of teleologically_k grounded norms. We can say that the norms in T 'serve' the end or value of knowing more: they are the norms subjects trying to know more ought to conform to. My thought is that we know that if a set of norms is teleologically grounded in some value v , then the reason why we ought or are permitted (and so on) to do as those norms say is that doing so helps to bring about v . So following the norms in T helps subjects get more knowledge. And if we think of T as exhaustive then if one is trying to know more, these are the norms that one ought to conform to.

The teleologist_k wants to say that the norms in E are those that a subject in pursuit of knowledge ought to conform to – those norms serve our goal of acquiring knowledge. Or given what I've just said about T, that E is a subset of T. I don't think that this can be right. I want to argue that some of the central

norms in E will not figure in T. And if we are thinking of T as the set of all norms that serve our knowledge acquisition ends, then the argument will show that those central norms aren't teleologically_k grounded.

Let's think about the sorts of norms in E then. Here is a natural place to start.

Weak Evidentialist Norm If one's total evidence clearly and decisively supports p at t , then believing p at t based on that evidence is permissible.

This norm just says that if a subject has decisive evidence for p and that this is clear to the subject, then the subject is permitted to follow that evidence and believe p . Of course, it's hard to say just what kind of evidential support counts as clear and decisive, but I think any intuitive way of fleshing it out will do. I take it that most will be happy to think of something like this sort of weak evidentialist norm as in E. Here are a few other norms that I think we can take to be in E.

Weak Reliabilist Norm Forming a belief p by way of a reliable belief-forming method is permissible.

Weak Knowledge Norm for Believing If one's believing p at t results in one's knowing p , then one's believing p at t is permissible.

Or relatedly,

Weak Knowledge Norm For every S , p and t , S 's coming to know p at t is permissible.

The sort of permissibility at work in all four of these norms is epistemic. That believing is epistemically permissible in these sorts of circumstances does not mean it's permissible all things considered. It just means that from some purely epistemic perspective believing under these conditions is permitted.

The last two knowledge-centric norms might be slightly less familiar so let me say something brief about them. My thought is that when a subject believes p and thereby comes to know p , they've done something that is at least epistemically permissible on our current way of thinking about normative epistemology. If you're in the position to know p and you come to know what you're in the position to, you've done something epistemically fine (at least). So these two norms

are just picking up on a fairly uncontroversial thought that if you've managed to come to know then you've done something epistemically right.

Plausibly, there are stronger versions of all four of these norms that say that the subject's believing is not just permissible, but justified, warranted, virtuous, rational or even required. For our purposes we only need to focus on these weaker norms.

Are norms like these going to be in T ? We've seen that there's at least a *prima facie* case that they will be: for instance, plausibly the reason you're always permitted to come to know is that knowing is valuable or valued. On further reflection though I don't think that subjects aiming to know more should conform to these norms in all cases and so I don't think we'll find them in T .

Rather than continue to discuss all of these norms here I'm going to stay focused on the Weak Knowledge Norm (WKN). This is simply for ease of exposition since it will be easy to see that most everything I say about WKN in what's to come applies just as well to the other three norms. So why isn't WKN in T ?

My thought is that T is the (exhaustive) set of norms that a subject trying to know more ought to conform to. If WKN is in T , then it too has to be a norm that a subject pursuing the aim of knowing more ought to conform to. This is just what the teleologist_k claims. But what does it mean to say that one should always conform to a norm like WKN? In particular, WKN is a permissibility norm and so conforming to it isn't a matter of always doing what it demands. What I think we should say in this case is this. Say that some norm n says that φ -ing is permissible in circumstances c . If n is in T then it should never turn out impermissible, by the lights of T , for a subject in pursuit of knowledge (in the relevant sense) to φ in c . Perhaps another way to think of it is this: when a requirement norm serves some end we can say that a subject in pursuit of that end ought always to do what the norm demands. And when a permissibility norm serves some end we can say that a subject in pursuit of that end is always permitted to do what the norm allows. This is to say that a subject conforming to the norms in T should always be permitted to do what n allows if n is in T .¹¹

¹¹Crucially, n here must be a 'pure permissibility' norm – one that doesn't itself contain certain kinds of limitations or provisos, e.g., *ceteris paribus* clauses. I take it that this is exactly how we should be thinking of WKN and the other norms in E . Those norms say that whenever one is in the relevant epistemic circumstances or conditions then believing is epistemically permissible. They don't say that that coming to know or believing on the evidence (etc.) are epistemically permissible except for cases in which other things aren't equal. Again, these are purely epistemic norms which make very general claims about the epistemic permissibility of believing in certain

I want to argue that a subject trying to know more is not always permitted to do what WKN permits – that subject is not always permitted to come to know what they're in the position to. And importantly, the cases in which subjects in pursuit of knowledge are not permitted to come to know are not unusual or exceptional. I want to argue that thinking about absolutely typical cases will get us to the conclusion that subjects trying to know more will often not be permitted to come to know.

Before we get to my main worry, a variant on a sort of case that's been in the literature recently (again) already makes trouble for the thought that WKN is in T. This is a sort of case in which one is in the position to know p but in which failing to come to know p will lead to a gain in knowledge at a later time. In these sorts of cases it seems as though a subject pursuing their knowledge acquisition goal won't be permitted to believe or come to know p at the relevant time.¹²

For instance, say S is the boss at Acme Inc. He's on the verge of a great discovery that is likely to change his field and lead to an explosion of very important knowledge. S is an awful boss and he treats his employees terribly, shouting at them regularly, paying them shabbily, offering them dismal working conditions and no health insurance. Each night S reflects on his day and how he performed. If he comes to think that his employees don't love him, he'll stay up all night fretting about his behaviour and worrying about whether he's a bad person and as a result will lose focus in the days that follow. If that happens the project will not succeed and he'll end up knowing much less than he otherwise would have. All the evidence he has makes very clear that he's despised.

On Monday night when S is contemplating whether he's beloved as a boss, what should he believe? Plausibly, from the perspective of the goal of knowledge acquisition, he should resist coming to know something he's in the position to. That's to say, it isn't permissible for him to come to know something he's in the position to come to know. S is a subject in pursuit of knowledge. If WKN is in T then it should always be permissible for S to do what WKN permits. But then this Acme case looks like a counterexample: it's a case in which a subject trying to know more is not permitted to do what WKN permits and so seems to show

kinds of circumstances. The norms in E that we're thinking about here are all pure permissibility norms in this sense.

¹²These sorts of cases can be found in a number of spots in the epistemology literature. They go back at least to Roderick Firth in e.g., Firth (1981).

that WKN is not in T and so not teleologically_k grounded.¹³

These sorts of cases can feel rarefied. If these cases are just at the far fringes then we might think that there are ways to dismiss their force as counterexamples. Perhaps with a little bit of tinkering we can fix our statement of the norms in E or think more about what is and is not in T to avoid these strange cases, leaving it open that norms like WKN are in T. I'm sympathetic to the thought that it can feel as though there is some wiggle room given just these sorts of cases. But I don't think that we need to turn to strange cases at all to see why E is not a subset of T. A teleological_k account of our epistemic norms doesn't just fail at the edges but in absolutely typical or normal cases (and all the time). Its basic structure makes it that WKN is not in T (and same goes for the other norms mentioned earlier).

To see this, let's start here. One sort of thing that the Acme case and others like it bring out is that once we see subjects as in pursuit of some valuable or valued end there is pressure to think of the relevant subjects as having to maximize or optimize with respect to that end. A subject who is trying to acquire knowledge ought to get as much of it as they can going forward; anything less seems irrational. If one is trying to acquire money, and one is faced with a choice between acquiring one dollar, ten dollars and one hundred dollars, one opts for the latter on pain of irrationality (other things equal, of course). This is to say that there is pressure to say that there is a norm in T that demands we maximize or optimize with respect to our knowledge acquisition. This sort of norm will say that subjects should always come to know as much as they can. In many cases this will require that subjects come to believe some *p* or *q*, but perhaps not always – as we've seen in the Acme case it seems to demand that our boss refrain from believing. Let's assume that this sort of demand to maximize is in T. Later I'll revisit this assumption and show that much of what's argued here could go through without it.¹⁴

¹³One might try to argue that these cases don't get us quite this far. Perhaps they show that there's a norm in T that says that when not believing *p* will ultimately lead to more knowledge acquisition than believing *p*, don't believe *p*. If this norm were in T alongside WKN this would leave T inconsistent, which is problematic. Moreover, these sorts of counterexamples seem to me to work in the direct way I'm claimed that they do. If some norm *n* is in T then it should always be permissible for someone in pursuit of knowledge to do what *n* permits. But these sorts of examples show that it won't always be permissible for subjects in pursuit of knowledge to do what some of our traditional epistemic norms permit.

¹⁴I hope that it's clear why the teleologist is pushed towards demanding the relevant kind of value maximization. In general once we're thinking of subjects as in pursuit of something of

One question that comes out of Acme-like cases has to do with whether the demand to maximize that's in T is a demand to come to know the most one can right away, or sometime further down the road. Let's say that 'judging' is the act of coming to believe. We can think of the demand to maximize that's in T as follows: it demands that we always make the judgment that leads to our having at least as much knowledge as any other available judgment would leave us with. But now take the Acme case. Our boss makes a judgment that has as its upshot a false and unreasonable belief. He could have instead made a judgment that had his knowing something as its upshot. So from the perspective of the immediate upshot of his judgment, our boss has not maximized with respect to his knowledge acquisition. Let's say that our boss then fails to 'immediately maximize'. Of course, what's at the heart of the case is that failing to immediately maximize with respect to one's knowledge acquisition is compatible with one's having maximized relative to some other, later time. If we look to some time, *t*, further down the road for our boss, then we can say that his false and unreasonable judgment lead (in some sense) to his knowing much more at *t* than he would have had he immediately maximized. So should we read the demand to maximize that's in T as a demand to immediately maximize or not? Some have argued that the demand to maximize that is relevant to epistemology is a demand to immediately maximize.¹⁵ If that's right then Acme-like cases seem to lose their force as counterexamples to the claim that E is a subset of T.

We don't need to take a subject's future prospects into account to see why WKN won't be in T though. Assume the demand to maximize that's in T is a demand to immediately maximize. All we need to focus on is this demand and its relation to the central norms in E to see why WKN won't be in T. To make sense of this demand in more detail we need some way of measuring how much knowledge one gets with each judgment. This is obviously no easy matter and I'll try to avoid the hardest bits of it by keeping to fairly clear cases.

I assume that in a typical case (world and time) a subject is in the position

value or that they value then it's natural to think that they should get as much of it as they can (other things equal). I think this is true for the epistemic instrumentalist as well as the epistemic consequentialist. For the instrumentalist I think the pressure to maximize with respect to one's knowledge acquisition comes in part from the demand to take the best means to one's end. If one's end is knowing more in general, then the best means to that end at any given time is knowing as much as one can.

¹⁵For instance, recall Foley's claim in footnote 9 – it's about "now" believing what's true. For some discussion of these sorts of temporal issues and their relation to epistemic consequentialism in particular see Percival (2002).

to make a number of different judgments. Some of these will be irrational or otherwise epistemically inappropriate but many will be grounded in evidence and many will be what we can call ‘k-judgments’. K-judging is the process or act of coming to know. K-judging is a kind of judging, the kind that results in one’s knowing something. Let’s focus on the subset of cases in which subjects are in the position to make some k-judgment. I take it that these too are typical cases, that we are typically in the position to come to know something or other. Let’s call these ‘k-cases’. In a typical k-case I think there’s quite a lot that subjects can come to know. For instance, a great deal of information is typically available in visual perception and so the average perceptual scene presents the subject with a great deal of potential knowledge – when we perceive an even moderately complex scene there is a lot we can come to know. In addition, a typical subject can easily add to their store of knowledge at a typical time by exploiting their evidence – by drawing out the deductive consequences of those things they already know, and by drawing other sorts of inductive conclusions from their bodies of total evidence. So there’s a lot a typical subject can come to know at a typical time.

Here’s a way to think about a typical k-case then: for each subject in a k-case there is a set of knowable propositions, Ψ_i . The elements of Ψ_i though can plausibly be at least partially ordered in terms of how much knowledge each yields when known. But just this thought makes serious trouble for the teleologist_k, I think. Say some subject is faced with a range of k-judgments they can make, some of which will result in more knowledge than others. Say that subject comes to know less than they could have. Given the demand to maximize in T what this subject did will have been impermissible.¹⁶ But we’ve stipulated that what they did was come to know some propositions. So from the perspective of T it looks as though coming to know those propositions was impermissible. The rest of this paper is going to flesh out and defend this line of thought.

To start drawing it out let’s focus on a very straightforward way of ordering the elements of Ψ_i .¹⁷ Say that the elements of Ψ_i are ordered roughly accord-

¹⁶Here and throughout I assume that ‘ought’ and ‘permit’ are duals, i.e., that one ought to φ iff one is not permitted not to φ . Alternatively one could take the demand to maximize as a prohibition against not maximizing and largely avoid the duality assumption.

¹⁷While I’m going to stay focused on this way of thinking about how much knowledge is available in a k-case, I think the sort of argument I am making here could be made if one favoured some other measure instead.

ing to their logical strength. We can say that all atomic propositions in these sets yield the same amount of knowledge when known. Conjunctions of those yield more knowledge when known; conjunctions of those conjunctions yield even more knowledge when known; and so on. Then we can talk about the logically strongest propositions in Ψ_i . In some cases there will be just one: the conjunction of all of the atomic propositions, although I assume that often that conjunction will be too complex to know. When that happens some weaker conjunction may be the strongest or we may have a few strongest knowable propositions in a case. In each Ψ_i let's say that α_i picks out one of these logically strongest propositions. A demand to maximize one's knowledge acquisition at the time of judgment will demand that the subject come to know α_i or something as strong as α_i at the time of judgment. But this will mean that in each of these cases coming to know anything strictly weaker than α_i (and just that) is impermissible.¹⁸

There is nothing special about k-cases. That's to say, I think that knowledge is typically available to us. Moreover, a typical k-case is one in which we can know a number of different things and these can be compared according to how much knowledge they yield when known. But this means that if there's a demand to maximize in play for the teleologist_k, then typical cases are ones in which we can come to know some propositions but what we did was epistemically impermissible.

To give a concrete example, imagine you are reading in your office. You wonder where your lucky pencil is and look around. You look to your pen jar and see it there coming to know that your lucky pencil is in your office (p). But at the time at which you glance over to your pen jar (t) you could have come to know more than p , e.g., that your lucky pencil is in your office and so are at least four pens or that your lucky pencil is in your office and so are at least four pens and so is a blue pen. Say α in this case is some conjunction that entails p . This will mean that just coming to know p (coming to know p and nothing else) at t is impermissible given that more knowledge is easily available at t . I don't think there is anything rarefied about this case, quite the reverse I think this is an absolutely typical case. If there's a demand to maximize then it will typically

¹⁸The 'just that' here is important. For now the only claim I'm making is that coming to know any of these strictly weaker propositions without also coming to know what one is required to come to know in these cases is impermissible. I'm going to say more about some of the subtleties here in the next section.

turn out that we can make some epistemically impeccable judgments – ones that result in knowledge, are based in the evidence, are the outputs of reliable methods – but nonetheless what we’ve done is epistemically impermissible.

Here’s another sort of example. Take a k -case in which a subject is presented with all sorts of visual information. Say they’re on the train heading to work, or just at home doing some cooking. In cases like these there is a great deal of knowledge subjects can acquire by coming to know what’s visually presented to them. Moreover, in cases like these subjects can also typically add to their store of knowledge by doing something other than taking in that visual information. In particular, I take it that at most any time a typical subject can add to their store of knowledge by drawing inferences from the bodies of knowledge they already possess. So our cook and train rider can add to their knowledge at the relevant times by drawing inferences from things they already know that might be completely unrelated to what’s presented to them visually, e.g., how many hours of sleep they got last week or whether any of the students in their class have names starting with ‘x’. But it’s not clear to me that these subjects can do all of these things at once, i.e., take up the visual information and draw these inferences. If they cannot then in cases like this – which, again, I think are absolutely typical cases – the logically strongest propositions in the relevant Ψ s won’t be a conjunction of all of the knowable propositions at the time. Say the strongest proposition in each of these Ψ s is some p that can be known perceptually. Then coming to know how many hours of sleep they had or facts about their students’ names is not permissible according to the teleologist $_k$.

These examples are extremely easily multiplied. They don’t rely on subjects finding themselves in strange or unfamiliar circumstances, rather they are just descriptions of the sorts of cases we regularly find ourselves in and the normative verdicts the teleologist $_k$ will have to issue about them.

The upshot here is trouble for the thought that our familiar epistemic norms are teleologically $_k$ grounded. If we think that epistemic subjects are or ought to be trying to know more, then plausibly they will be expected to act so as to maximize with every judgment. This has the result that coming to know some propositions can very easily be impermissible. This means that WKN is not in \mathbb{T} . Recall, the condition WKN had to meet to be in \mathbb{T} was that it’s always permissible for a subject in pursuit of knowledge to do what it permits. But now we can see that once a demand to maximize is in \mathbb{T} , it won’t always be permissible for that

sort of subject to do what WKN permits; in fact, it very often will not be. So WKN is not teleologically_k grounded.

It should also be clear that what I've said here applies to the other norms articulated earlier as well, e.g., the evidentialist and reliabilist permissibility norms. In all the cases I've described what's impermissible is making some epistemically impeccable judgments – these result in knowledge, are the output of reliable methods, and are based in clear and decisive evidence. So I think the arguments here show that the rest of the standard epistemic norms I've specified are also not in T and so E is really not a subset of T.

If E is not a subset of T then we should also say that at least some of the norms in E are not teleologically_k grounded – they are not the norms subjects trying to know more ought to conform to; they don't serve our end of knowing more. This isn't to say that subjects in pursuit of knowledge shouldn't ever follow their evidence or use reliable methods or come to know, rather what we've seen is that once a demand to maximize is in play, there will be significant constraints on which evidence subjects should follow and what they should come to know. It's understandable, I think, to want some sort of explanation as to why we should conform to our central epistemic norms. I've argued that one natural explanation fails.¹⁹

4 Objections and replies

In this section I want to raise a few objections and discuss a few replies.

4.1 Just coming to know *p*

I've argued that teleological_k epistemology has the result that in a typical case there are all sorts of propositions that we are not permitted to come to know. But if we look to one example I gave, it might look as though that's not quite what we should say. I said that just coming to know that your lucky pencil is in your office (*p*) was impermissible according to the teleologist_k. So what's impermissible in this case is coming to know *p* without also coming to know the other things the teleologist_k demands you know in that case. This is to say that coming to know

¹⁹I ran this argument assuming that the demand to maximize that's in T was a demand to immediately maximize, but I hope that it's clear that the argument certainly does not rely on restricting the demand to maximize in this way.

p alongside everything else you're required to come to know is permissible, but coming to know it without coming to know all of the rest is not. What should we say about the (bare) permissibility or impermissibility of coming to know p then? Fully answering this question is too big a task for this paper, I think. I do want to make a few key points though. Crucially, there does seem to me to be something wrong with claiming without qualification that in this case judging or coming to know p is impermissible. With that in mind one might wonder whether (a) there are any cases in which it is impermissible without qualification or tout court to come to know some p at t for the teleologist $_k$; and (b) what impact cases like the lucky pencil case should have on the conclusion that norms like WKN are not in T .

On (a), I think it's easy enough to find cases in which the teleologist $_k$ should say that coming to know p at t is impermissible tout court. Recall, in each k -case we have a set of knowable propositions for each subject, Ψ_i . Let's continue to assume that the ordering that matters for our purposes is logical strength. I don't think it will be the case that the strongest propositions in Ψ_i will be conjunctions of all of the (atomic) propositions in Ψ_i , or even that for each Ψ_i there will be one logically strongest proposition. My thought here is that in many k -cases conjunctions of all the (atomic) propositions in Ψ_i are too complex for typical subjects to grasp and so won't be in Ψ_i . Instead some strictly weaker propositions will be the strongest or among the strongest. In cases like this there may be propositions in Ψ_i that it's impermissible to come to know, tout court.

For instance, imagine a subject with the following Ψ at t : $\{p, q, r, (q \wedge r)\}$. Here $(q \wedge r)$ is the strongest proposition and so the one that the relevant subject ought to come to know at t according to the teleologist $_k$. In a case like this I think the teleologist $_k$ should say that coming to know p at t is impermissible tout court and not just that coming to know p without coming to know some other things is. And we've already seen cases with this structure in the last section. The cases of the cook and the train rider in which subjects were in the position to both come to know some things perceptually and some other things by drawing inferences from already existing knowledge (but couldn't come to know all of that at once) were cases like this.²⁰

²⁰I've been largely assuming a picture according to which subjects are choosing between making different single judgments at a time. But I take it that it's possible to make multiple judgments at the same time. Allowing for subjects to make multiple judgments at once though won't change the conclusions here. For instance, one might try to argue that in the case just described the

Moreover, and on (b), I think even cases like the lucky pencil case make trouble for the teleologist_k. Here is my thought. We know that given the demand to maximize that's in T the teleologist_k must say that in many cases coming to know p without also coming to know more is impermissible. Take a case then in which a subject comes to know p and nothing else but ought to have come to know more. What that subject did was impermissible by the lights of T.

The question then is what the phrase 'what the subject did' picks out. On the one hand what the subject did was come to know p . On the other, one might think that what they did was more complex than that, perhaps involving coming to know p as a part and then failing to know some other things as a part. I think that these are interesting and important issues, but they aren't ones that I can fully engage with here. Instead, let me simply propose one way of thinking about these cases.

There are many ways in which one may come to know p : one can come to know p by just coming to know p , or by coming to know $(p \wedge q)$, or by coming to know p alongside all sorts of other things. WKN says that it's always permissible to come to know and I take it that this extends to all of the sorts of ways of coming to know just mentioned. From the perspective of epistemology alone, if you form a belief and that belief amounts to knowledge then the judgment was epistemically permissible, no matter whether you came to know p alone or also came to know more. But it looks as though many ways of coming to know p can be impermissible for the teleologist_k, e.g., just coming to know p in the sorts of cases we've been thinking about.

Here's another way to think of this point. If you come to know p then we can assume that the resulting belief p is epistemically justified. And this applies to any of the sorts of ways of coming to know just mentioned. But the following general principle seems extremely plausible: if some belief is epistemically justified then it was epistemically permissible for you to have formed that belief. So from the perspective of E any way of coming to know p is permissible, but from the perspective of T, many ways of coming to know p are not.

subject is permitted to come to know p at t so long as they also come to know $(q \wedge r)$ at the same time. But part of the point of the class of cases being described here is that there is a limit on how much of the "available" information the relevant subjects can come to know at once. In these sorts of cases coming to know some things at the relevant time will make it that one cannot come to know some other things at that time. This isn't because we cannot in general make multiple judgments at a time though.

4.2 Other forms of teleological_k (and teleological) epistemology

I've said that the teleologist thinks that the explanation of why we ought to conform to our central epistemic norms goes via the relation those norms bear to epistemic values. It's doxastic or epistemic states that are valuable or valued and I've assumed that the teleologist wants to ground our epistemic norms in the value of being in the states of value, e.g., of having knowledge. But one might argue that a teleological account of our norms needn't take that form. For instance, perhaps rather than claiming that our epistemic norms are the norms subjects trying to know more ought to conform to, an epistemic teleologist can claim that our epistemic norms are the norms subjects trying to avoid beliefs that don't amount to knowledge ought to conform to. On a sort of view like this perhaps epistemic subjects are subjects aiming to know p whenever they judge p (but not necessarily aiming to know more or aiming to know in some general sense). And our epistemic norms are norms for these sorts of subjects. This might still be a way of thinking of our epistemic norms as teleologically_k grounded.²¹ And it might avoid some of the worries I've raised here which stem largely from the thought that valuing knowledge involves trying to get more of it. I can imagine other sorts of broadly teleological_k (and teleological) treatments of our epistemic norms as well that might still be on the table given what I've said here. I don't mean to be saying that views like this couldn't count as teleological_k and so I don't mean to be saying that there is no room at all for a teleological_k treatment of our epistemic norms. That said, I think that the most natural teleological_k views will take it that epistemic norms are norms for subjects who are actually trying to know things. The argument here makes trouble for those sorts of views.

I also argued that once we think of epistemic subjects as subjects who are or should be trying to know more, a demand to maximize comes naturally. That said, I think much of what I argue here could retain its force even if one insisted

²¹Of course it remains to be seen whether a picture like this could ground our epistemic norms. My suspicion is that the picture of epistemology that results, although very interesting, would be at least somewhat revisionary. For instance, while perhaps it does go some way towards grounding the sorts of permissibility norms discussed in this paper, on this sort of picture it looks as though epistemology, on its own, doesn't really give us reasons to form or have beliefs. Rather it just gives us advice about which beliefs to form or have on some subject matter if we're going to form or have any. It's interesting to think about how well this sort of picture comports with the general thought that knowledge is valuable as well. We'll have to save that for another time, unfortunately.

that some weaker sorts of value promotion principles were in T in the place of a demand to maximize. For instance, say T contained only a demand to satisfice with respect to one's knowledge acquisition rather than maximize. That norm won't say that subjects ought to come to know as much as they can with every judgment but will nonetheless say that there is some limit on how little they are permitted to come to know; anything below that threshold will be impermissible. And this leaves the satisficing teleologist_k in much the same position as the maximizer: the relevant familiar epistemic norms we've been thinking about still won't feature in this sort of satisficing teleological_k epistemology. More generally, so long as T contains a norm that maintains a permissibility threshold for knowledge acquisition such that subjects can come to know p but ought to have come to know more, the teleologist_k is in much the same position.

And finally, as I said at the start of the paper, even though I've focused on teleological_k epistemology, I think other forms of teleological epistemology – ones in the which the valuable state is something other than knowledge, or ones in which knowledge is one of many epistemic states of value – are at risk in the same sorts of ways. If the relevant values, e.g., true belief, understanding, are states we're trying to get into, then I think we can run analogous arguments to the one I ran here.

4.3 Important knowledge

Perhaps there is a different way for the teleologist_k to avoid the idea that we ought to come to know as much as we can at every turn. My assumption at the outset was that knowledge in general is valuable – all knowledge. But that claim has been denied with some arguing that only some knowledge, only “important” knowledge, is valuable.²² One might try to argue then that the demand to maximize that's in T is only a demand to maximize with respect to important knowledge. This might appear to help the teleologist_k here. Going back to the lucky pencil case, I argued that just coming to know p (that your pencil is in the office) was impermissible at t because there was more you could have come to know at that time. But if all that matters to you is where your lucky pencil is, then perhaps those stronger propositions are unimportant leaving p as the strongest knowable important proposition. In that case just coming to know p might turn out perfectly permissible for this importance-focused teleologist_k.

²²See Grimm (2009) for some of this debate.

I don't see much help in the end here though. Let's imagine we have some way to distinguish important from unimportant knowledge, even roughly. Now say that only important knowledge is valuable or valued, and so that only important knowledge ought to be acquired (and we'll have to assume that acquiring unimportant knowledge is at least permissible lest there be plenty we're not permitted to come to know according to T). So long as there are cases in which multiple important propositions are knowable, we can still make sense of ordering these in terms of their knowledge yield (and importance now?). And so long as there's a demand to maximize or satisfice in place this teleologist_k is in much the same position as before.²³

4.4 Acts and states

Finally, I've argued that the relevant norms in E are not going to be in T since there will be a norm in T demanding that subjects maximize with respect to their knowledge acquisition and given that norm it is false (from the perspective of T) that subjects are always permitted to come to know (and so on). But the demand to maximize is a demand to do something: it's a demand to make certain kinds of judgments.

One might try to argue that traditional epistemology consists in "state-focused" norms only: that all of its norms tell us how to be rather than what to do and given this that the argument here missed its mark since it only shows that certain kinds of judgments are impermissible according to the teleologist_k leaving it open that the resulting beliefs be permissible. Let me make a few points here.

First, while contemporary normative epistemology certainly does care about the evaluation of doxastic and epistemic states, traditional epistemic norms also bind us with respect to our judgment making. Take a subject who has no view about whether *p* is true. They uncover decisive evidence that it is. I take it that we don't want to say that epistemology is silent about what this subject is permitted (or even required) to do. This subject is epistemically permitted (or even required) to judge *p* – to come to know what they're in the position to come

²³Moreover, as I said earlier in the paper in discussing global versus local forms of epistemic instrumentalism, the epistemic norms the teleologist_k wants to explain or ground are very general: they permit or require or justify believing things on any subject matter or topic. How could a restricted goal like acquiring knowledge on only some subject matters ground (say) a permission to believe whatever one's evidence supported?

to know.²⁴ Our epistemology is also committed to the permissibility of certain judgments and those judgment norms will not be teleologically_k grounded.

Moreover, there is a case to be made for the teleologist_k's having to move from the claim that some judgment is impermissible to the claim that the resulting belief is. This is because the sort of teleologist we've been thinking about takes it that the normative statuses of the relevant acts of judgment are derived from the normative statuses of the states that result from those acts. Our teleologist_k proceeds by ranking outcomes according to how valuable they are, and then deeming the acts that lead to those various outcomes as permissible or impermissible in large part based on the values of those outcomes. For instance, it's in virtue of the fact that judging *p* leads to the state of affairs in which knowing is maximized that judging *p* is permissible for the teleologist_k. This means that if judging *p* is permissible according to the teleologist_k that has to be because the resulting state (or state of affairs) is sufficiently good or valuable, and if judging *p* is impermissible that's because the state that results is not sufficiently valuable. So if the teleologist_k says that some judgment is impermissible, they are also saying something about the disvalue of the state that results from that judgment.

Of course this isn't yet the claim that it's impermissible to be in the relevant states, but I think if the teleologist_k is going to make any deontic pronouncements about the states, that's exactly the one they ought to make. And I don't see any special reason to think that they cannot make some deontic pronouncements about these states. That the normative structure of the teleologist's position is "state first" typically means that they start with some axiological evaluations of the states and from there arrive at some deontological evaluations of the acts leading to the states. In epistemology though we want some deontic evaluations of the states as well, not just the acts. How should those be distributed for the teleologist_k? Surely the most natural way is in line with those of the acts that lead to the states. Any other arrangement would need some explicit motivation. If this is right though then the most natural thing for the teleologist_k to say is that if some judgment is impermissible then the resulting belief state is as well.

But if this is right then the arguments in this paper also show that the state-focused analogues of the central epistemic norms we've been thinking about are

²⁴In fact see Kolodny (2005) for the thought that all rational requirements are "process requirements".

also not in T , e.g., norms that say that having a belief based in the evidence or having a belief that's the product of a reliable belief-forming method or knowing are always epistemically permissible.

5 Concluding remarks

Could the teleologist_k argue that norms in E are helpful rules of thumb for subjects trying to acquire knowledge? Given what's been argued here, this doesn't seem plausible. So long as there is pressure to think of subjects in pursuit of knowledge as having to maximize (or even something less demanding), then the norms we've looked at simply don't look like good rules of thumb for subjects trying to know more. In a typical case there will be much subjects are not permitted to come to know.

We like to think that we are subjects in pursuit of valuable epistemic ends. And even that epistemology provides us with rules we should conform to in pursuit of these goals. I've argued though that our familiar epistemic norms aren't that. They aren't the norms a subject trying to know more (or have more true beliefs, etc.) should conform to; they aren't teleologically grounded, at least not in the natural sense I've been focused on. This doesn't spell the end of this sort of teleological epistemology though, only the end of this sort of teleological treatment of our central epistemic norms. One might see in some of the arguments here a case for embracing some other epistemic norms. If you think that (say) knowledge is valuable and that our epistemic norms should be the norms subjects in pursuit of that good should conform to, then I hope to have shown that you'll need to embrace some new epistemic norms and to have given at least a hint at what those norms might look like.²⁵

²⁵Thanks very much to Dave Chalmers, Charles Côté-Bouchard, Harvey Lederman, Jake Nebel, Amia Srinivasan, and an anonymous referee.

References

- Ahlstrom-Vij, K. and Dunn, J. (2014). A Defence of Epistemic Consequentialism. *Philosophical Quarterly*, 64(257):541–551. 4
- Berker, S. (2013a). Epistemic Teleology and the Separateness of Propositions. *Philosophical Review*, 122(3):337–393. 4
- Berker, S. (2013b). The Rejection of Epistemic Consequentialism. *Philosophical Issues*, 23(1):363–387. 4
- Cowie, C. (2014). In Defence of Instrumentalism About Epistemic Normativity. *Synthese*, 191(16):4003–4017. 5
- David, M. (2001). Truth as the Epistemic Goal. In Steup, M., editor, *Knowledge, Truth, and Duty*, pages 151–169. New York: Oxford University Press. 2
- Feldman, R. (2000). The Ethics of Belief. *Philosophy and Phenomenological Research*, 60:667–695. 5
- Firth, R. (1981). Epistemic Merit, Intrinsic and Instrumental. *Proceedings and Addresses of the American Philosophical Association*, 55(1):5–23. 9
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Harvard University Press. 5
- Grimm, S. (2008). Epistemic Goals and Epistemic Values. *Philosophy and Phenomenological Research*, 77(3):725–744. 5
- Grimm, S. (2009). Epistemic Normativity. In Haddock, A., Millar, A., and Pritchard, D., editors, *Epistemic Value*. Oxford: Oxford University Press. 19
- James, W. (1896). *The Will to Believe*. New York: Longmans, Green and Co. 2
- Joyce, J. M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65(4):575–603. 2
- Kelly, T. (2003). Epistemic Rationality as Instrumental Rationality: A Critique. *Philosophy and Phenomenological Research*, 66(3):612–640. 5
- Kolodny, N. (2005). Why Be Rational? *Mind*, 114(455):509–563. 21

- Kornblith, H. (1993). Epistemic Normativity. *Synthese*, 94(3):357–376. 5
- Leite, A. (2007). Epistemic Instrumentalism and Reasons for Belief: A Reply to Tom Kelly’s “Epistemic Rationality as Instrumental Rationality: A Critique”. *Philosophy and Phenomenological Research*, 75(2):456–464. 5
- Percival, P. (2002). Epistemic Consequentialism. *Aristotelian Society Supplementary Volume*, 76(1):121–151. 11
- Steglich-Petersen, A. (2011). How to Be a Teleologist About Epistemic Reasons. In Resiner, A. and Steglich-Peterson, A., editors, *Reasons for Belief*. Cambridge University Press. 5
- Stich, S. (1990). *The Fragmentation of Reason*. MIT Press. 5
- Street, S. (2009). Evolution and the Normativity of Epistemic Reasons. *Canadian Journal of Philosophy*, 39:213–248. 5
- Talbot, B. (2014). Truth Promoting Non-Evidential Reasons for Belief. *Philosophical Studies*, 168(3):599–618. 1, 5