

November 2021. Draft version. For *Contemporary Debates in Epistemology, 3<sup>rd</sup> Edition*, (eds.) E. Sosa, M. Steup, J. Turri, & B. Roeber, (Wiley-Blackwell, *forthcoming*).

## Epistemic Normativity is not Independent of our Goals

J. Adam Carter  
*University of Glasgow*  
adam.carter@glasgow.ac.uk

According to epistemic instrumentalists, epistemic facts (including about what we ought to believe, or about what is an (epistemic, normative) reason to believe what) obtain at least partly in virtue of our goals (or aims, ends, intentions, desires, etc.).<sup>1</sup> In this respect, epistemic normativity is *dependent* on our goals. *Anti-instrumentalists* deny this: they maintain that such facts don't obtain even partly in virtue of our goals (broadly construed, to include aims, ends, intentions, desires, etc.).

In his contribution to this volume, Alex Worsnip ([Forthcoming](#)) defends anti-instrumentalism. While I reject many specific ways to defend epistemic instrumentalism – including some of the very same species of instrumentalism that Worsnip challenges – I think we should nonetheless reject anti-instrumentalism. Strictly speaking, it's false that epistemic normativity is entirely independent of our goals (where 'goals' is construed broadly enough to include *aims*). In fact – on the type of view I favour, *telic virtue epistemology* – epistemic normativity is *inextricably* aim-dependent – viz., relative to the aims (both intentional and teleological) of thinkers in trying to answer

<sup>1</sup>I borrow this specific characterisation of epistemic instrumentalism from Worsnip ([Forthcoming](#), X).

questions correctly (or knowledgeably).<sup>2</sup> That said, the sense in which telic virtue epistemology maintains that epistemic normativity is aim dependent is one that easily ‘falls between the cracks’ in typical debates between epistemic instrumentalists and anti-instrumentalists.

§1 does a bit of scene setting to explain *why* this is so – and in the course of doing so, I’ll clarify the distinction between *prescriptive* and *evaluative* epistemic norms; and between *substantive* normativity and *non-substantive* normativity. Reference to these distinctions will help us see just how space opens up to resist epistemic anti-instrumentalism *without* signing up to *any* of the specific versions of instrumentalism that Worsnip explicitly takes issue with. §2 then lays out the *non-substantive, evaluative* epistemic normativity that is central to telic virtue epistemology, and in doing so, clarifies various ways in which epistemic normativity *is* aim dependent. §3 makes a few remarks in order to distinguish epistemic norms from wider, substantive zetetic norms, viz., norms that govern *inquiry*. §4 then revisits and sharpens of the key points of agreement between myself and Worsnip before registering where I take our dispute to really lie (while emphasising where it does not).

## 1. Some normative distinctions

a. An assassin’s shot, from 100 feet away, might be an *excellent* shot – released with great skill and succeeding in its intended aim through that skill. At the same time, the shot might very well be a moral abomination. It is to the assassin’s *moral* discredit that the shot hit rather than missed; but, at the same time, the shot is *creditable* to the assassin, and in a way that it is not creditable to (say) any of the would-be assassins who missed. This is *telic* credit: attributable to the agent in a way that does not import any further axiological standing.<sup>3</sup>

To get at this idea from another angle, consider that – from the fact that the assassin’s shot is “excellent” and that, when it hits its target, its doing so is creditable to the assassin – it doesn’t follow that the assassin *should* ever have taken the shot in the first place, or for that matter, that they ever had any (substantive) normative reason to take the shot.

b. What is going on here? Why, then, are we *positively* evaluating a shot that, we might suppose, shouldn’t have been taken, and which we may even suppose the shooter had no (substantive) normative reason to take?

<sup>2</sup>This is just a first approximation; the details will be filled out in §2.

<sup>3</sup>For discussion, see Sosa (2021, 18).

We are doing what we do when evaluating a knife as a *good knife* when it is sharp. The kind of norm implicit in this evaluation is an *evaluative norm*. Evaluative norms – unlike *prescriptive norms*, which prescribe conduct – *regulate what it takes for a token of a particular type of thing to be good or bad with regard to its type*, where the ‘goodness’ here is attributive in Geach’s (1956) sense – viz., the sense in which a sharp knife is a good knife, *qua* knife, regardless of whether it is good or bad simpliciter.<sup>4</sup>

When Stephen Curry makes a 30-foot 3-point shot attempt, we say that it is a *good shot* in this same sense. We are not saying it is good *simpliciter*<sup>5</sup>, just that it is good *as an instance of the kind of thing it is* – viz., *qua* attempt to put the ball in the basket.

Telic virtue epistemology of the sort I embrace holds that the evaluative normativity of beliefs – aimed intellectual shots<sup>6</sup> – is like the evaluative normativity of aimed basketball shots. We’ll unpack this analogy in some more detail in §2.

c. Let’s now contrast Stephen Curry’s shot inside the area with what is happening outside the arena: suppose a medic standing in the ticket line hears a man choking, food lodged in his windpipe, struggling for air. The medic disregards her tickets and runs to help; she performs the Heimlich maneuver to perfection, saving the man’s life.

True, the medic’s action here is good *as an instance* of the Heimlich maneuver. But we think the medic’s action in this case *also* has *substantive* moral worth; it is to the medic’s moral credit that they performed as they did. The success of the performance is not *merely* creditable to them in the way we credit an archer with a good shot, or for that matter, the way a *trainee* medic’s performance of the Heimlich maneuver (to perfection) is creditable to them when performed just for fun, in training.

d. Consider this reasoning. *Beliefs*, when known, are *substantively* good, not merely *qua* beliefs. They are good not *just* like Steph’s shot (which is a good token of its type, in that it attains the aim internal to the kind of attempt it is), but also like the medic’s performance of the Heimlich maneuver, which is of substantive worth.

<sup>4</sup>For a helpful overview of the prescriptive/evaluative norm distinction, with reference to attributive as opposed to predicative goodness, see McHugh (2012, 22); cf., Simion et al. (2016, 384–86).

<sup>5</sup>It *might* be good simpliciter, perhaps, if it leads to great happiness, etc. However, that the shot is good simpliciter is not communicated through the use of “good shot.”

<sup>6</sup>The sense of ‘aim’ here will be unpacked in §2.

An assumption that might lead to this reasoning is that getting to the truth, or knowing, is a *substantively* worthy aim, like saving a life, or keeping a promise.

This assumption, though is difficult to defend, when we consider the burden that would fall on one who hopes to defend the substantive worth truth *per se* – applicable to any and all truths. Even if the substantive moral worth of keeping a promise is applicable *to any and all* instances of promise keeping, the same does not seem so in the case of truths. (Think, for example, of counting grains of sand on a beach.)<sup>7</sup>

Here's Ernest Sosa (2021):

Truth is not *per se* substantively valuable. Many truths are entirely worthless in themselves (absent curiosity or some other reason for their pursuit). And the same goes for one's knowledge of a worthless truth. Such knowledge is also deprived of any value and of any deontic requirement that it be pursued by anyone (again, absent curiosity or other such reason that might favor its pursuit). (2021, 7).

Note that the idea that truth *per se* is not substantively valuable does *not* mean that your believing a substantively worthless truth is not in the market for being good *like Steph's shot is good*, when the attempt at getting it right you make when affirming the pointless truth hits its mark. It might after all, like Steph's shot, be a good token of its type, in that it attains (skilfully so) the aim internal to the kind of attempt it is.

e. Any basketball shot might be good in at least *three* different ways, *qua* the kind of attempt it is: it might go in, in which case it is *successful qua* that kind of attempt, even if the shot, incompetently heaved at the basket, goes through the net just by luck; regardless of whether it actually goes in, it might be shot with excellent form, in a manner that reliably (enough) succeeds, in which case it is good *in that it is competent*; and – better yet – it might be good in that it is *successful because competent* – viz., it might be *apt*.<sup>8</sup>

<sup>7</sup>For a different take on this point, see Kvanvig (2008).

<sup>8</sup>For a helpful presentation of this idea, see Sosa (2010); cf., for a similar presentation, Greco (2010, Ch. 5).

Success, competence, and aptness, are thus three *evaluative norms*, applicable to basketball shots as well as to other aimed attempts.<sup>9</sup> Because beliefs are aimed attempts (in a sense I'll unpack further in §2), all three evaluative norms are applicable to beliefs, where – on the view I favour – *apt belief* lines up with *knowledge* and *competent belief* lines up with *epistemically justified belief*.

If this kind of thinking is on the right track, then it looks like we've got – with telic virtue epistemology – a qualified but not insignificant kind of incompatibility with epistemic anti-instrumentalism. At least, if 'epistemic normativity', as it features in the anti-instrumentalist's core thesis, is interpreted widely enough to include not only substantive, prescriptive norms (such as norms about what we ought to believe on balance, and what is a reason to believe what), but also *non-substantive, evaluative norms*; and (ii) 'our goals' include telic aims, then (iii) from the core ideas at the heart of telic virtue epistemology, it follows that epistemic normativity depends on our goals.

*This* specific way of denying epistemic anti-instrumentalism, though, lines telic virtue epistemology up with a species of epistemic instrumentalism that is *significantly* different from the kind of instrumentalist theses that anti-instrumentalists typically have in their crosshairs.

To use just two examples, consider first Richard Foley's (1987) goal-dependent articulation of an epistemic norm on (substantively) rational belief:

It is epistemically rational for s at t to believe p = df. (1) **S has the goal at t of now believing truths and now not believing falsehoods**; (2) S has at t an argument for p that S would regard as sufficiently likely to be truth preserving were S to reflect at t upon what else S believes at t and upon S's deepest epistemic standards at t and (3) S at t would have no reason to be suspicious of the premises of the argument [...] were S to reflect upon what else S believes at t and upon S's deepest epistemic standards at t. (1987, 66–67).

Foley's position parts ways with anti-instrumentalism because it countenances a *substantive* epistemic norm (one governing rational belief) that is goal dependent. The anti-instrumentalist rejects this proposal by rejecting that clause (i) has a rightful place in an account of substantively rational belief.

<sup>9</sup>These three are not exhaustive – e.g., *full aptness* is a higher quality telic norm than aptness, though one we can set aside for our present purposes to simplify discussion. For discussions of full aptness, see, e.g., Sosa (2015, 2021) and Carter (2020).

Consider likewise Hilary Kornblith's (1993) evidentialist (substantive) norm on belief, on which you ought, on balance, to believe in accordance with the evidence.<sup>10</sup> What *grounds* this norm, for Kornblith, is that we have practical goals that are in fact best served by our having accurate representations of the world.<sup>11</sup> And so the substantive 'ought' that features in the prescription to believe in accordance with the evidence is grounded in our practical interests and goals.

The anti-instrumentalist rejects *this* by denying that what one *on balance* ought to believe (viz., a substantive norm on belief) depends on practical interests, and this, the anti-instrumentalist says, is so even if (as in Kornblith's case) the norm itself prescribes believing in accordance with evidence.

Note that *both* Foley and Kornblith embrace goal-dependent articulations of specific substantive epistemic norms, and in doing so, are committed to the view that *at least some substantive epistemic norms are goal dependent*.

It should now be apparent why the specific way in which the telic virtue epistemologist is committed to a (qualified) denial of anti-instrumentalism easily falls between the cracks *in debates between* anti-instrumentalists and *instrumentalists* like Foley and Kornblith. Debates between anti-instrumentalists and Foley- and Kornblith-style instrumentalists are debates about whether substantive epistemic norms are goal dependent. And on this point, the telic virtue epistemologist *trivially aligns with the anti-instrumentalist*, rather than with Foley- and Kornblith-style instrumentalists. For the telic virtue epistemologist who thinks, as Sosa puts it, that "[...] unlike ethics, epistemology focuses on objectives neutral in substantive normativity", there are no substantive epistemic goal-dependent norms *because there are no substantive epistemic norms*<sup>12</sup>; (and, not because, as the anti-instrumentalists like Worsnip hold, because there *are* substantive epistemic norms though none of these are goal dependent).

## 2. Telic virtue epistemology and the aim dependence of evaluative norms

a. Does the telic virtue epistemologist's analogy between beliefs and other aim-directed attempts (like basketball shots) break down?

<sup>10</sup>See Kornblith (1993, 370–71). Cf., for discussion, Sharadin (2018).

<sup>11</sup>This kind of view is open to what Worsnip calls the 'trade off' problem. See Worsnip (this volume, §4).

<sup>12</sup>This idea will be unpacked further in §3, when we'll see that the telic virtue epistemologist *can* countenance substantive norms of *intellectual ethics*, which fall outside epistemology proper, and which are norms governing how we put ourselves in a position to know, rather than, norms governing the constitutive attempts we make at knowing.

The basketball player *really aims* at the basket. An archer *really* aims the bow and arrow at the target. But, as Ralph Wedgwood (2002) puts it: “Beliefs are not little archers armed with little bows and arrows: they do not literally ‘aim’ at anything” (2002, 267).

Beliefs are not archers! But *persons* are, and persons are also believers. This is relevant here because the sense in which the virtue epistemologist holds that *belief* aims at truth is derivative<sup>13</sup> upon the more fundamental idea that the person *by virtue of* believing, aims at truth.<sup>14</sup>

Even so, *how* exactly should we understand the believer’s aiming? How if at all does this differ from the kind of goal-directed aiming we see in (substantive) instrumentalists like Foley and Kornblith?

b. An initial way to gain some clarity here will be to take the analogy between believing and basketball shots (and other aimed performances) *more* rather than *less* seriously; *both* are species of the more general phenomenon of an *attempt*, where *attempts* are (on telic virtue epistemology) the central phenomenon of evaluation. Attempts are of central focus in telic virtue epistemology because (i) telic normativity is the (evaluative) normativity of attempts *as* attempts; and (ii) on telic virtue epistemology, the normativity of *knowledge* is a special case of telic normativity (Sosa 2021, 22).<sup>15</sup>

That said, we can narrow down the idea of an attempt itself into two categories – *instrumental attempts* and *constitutive attempts* – to get a grip on how specifically *knowledge-apt* attempts are related to their aims (though we’ll develop this grip further in the §2.c).

Suppose a golfer (say, Phil Mickelson) is on the final hole of a tournament, and needs to sink a 15-foot-putt in order to force a playoff. He aims to make it, no doubt, though there are several different means we can expect him to implement in attempting to secure that objective.

First, he’ll *line up the putt*, by standing behind the ball, gauging the slope of the green, which helps him to assess how hard to hit it, and whether the putt should begin its course left or right of the hole. These *means* (viz., gauging the slope, etc.) are viewed by Phil as just preliminary means in the service of his objective, means by which he puts himself *in a position* to make the putt and

<sup>13</sup>Note that there is alternative and well-known route to the conclusion that ‘belief aims at truth’ that needn’t go through telic virtue epistemology – and which is pursued by proponents of *normativism about belief* (see, e.g., Shah 2003; Shah and Velleman 2005; cf., for criticism, Glüer and Wikforss 2009).

<sup>14</sup>This is an approximation that will be filled in in some more detail in §2(c).

<sup>15</sup>Recall: on this view, knowledge is type identical with *apt belief*.

thereby attain his objective.<sup>16</sup> His lining up the putt, and implementing these preliminary means, is an *instrumental* attempt, not a *constitutive* attempt, to make the putt.<sup>17</sup> Hitting the ball is *itself* a constitutive attempt to make the putt, not merely an instrumental attempt.<sup>18</sup> When Phil *actually hits the ball* (rather than just lining it up) he attempts to make it through means (e.g., taking the putter back and releasing it through the ball) *whereby he would succeed*, grounding his success.<sup>19</sup>

For a telic virtue epistemologist, the case of Phil and the case of an inquirer are analogical in that *lining up a putt* is to *hitting the putt* as *gathering evidence* is to *believing*.<sup>20</sup>

What beliefs, basketball shots, and putts have in common – all of which can be assessed for success, competence, and aptness – is that they are *constitutive attempts*; all of these attempts involve the taking of means that would *ground* the relevant success, rather than merely *put one in a position* to do so.<sup>21</sup>

What the forgoing suggests, then, is that knowledge apt attempts – viz., beliefs – are attempts that are aim-related *in the way that constitutive attempts more generally are aim-related*, as opposed to how other, merely instrumental attempts, are aim-related. This gives us at least a provisional grip on how knowledge-apt attempts are aim-related (and *a fortiori* how knowledge is aim- or goal-*dependent*), but it's only part of the story.

c. Even if we grant that Steph Curry's shot is a (constitutive) attempt to

<sup>16</sup>This is a variation on a distinction we find in Sosa (2021, 22).

<sup>17</sup>*Ibid.*, 22.

<sup>18</sup>When Phil does *that* he is also implementing means aimed at the objective (of making the putt); but in this case, the means he implements are not viewed as just preliminary (which would put him in a position to make the putt), but as actually *grounding* the success itself.

<sup>19</sup>For fuller discussion on these points, see Sosa (2021, Ch. 2).

<sup>20</sup>Whereas Phil wants to make the putt, an inquirer (into whether *p*) wants to *settle the question* of whether *p*. Phil instrumentally attempts to make the putt by lining it up and the inquirer instrumentally attempt to *find out whether p* by gathering evidence, weighing it carefully, etc. That is to say, by implementing means viewed as preliminary, viewed as means that would *put one in a position* to know whether *p*. Further developing this analogy, just as Phil (after whatever instrumental attempts are made) hits the ball and in doing so *constitutively attempts* to make the putt (taking means that would ground his success in the objective), you after (after whatever instrumental attempts are made) *form a belief* as to whether *p*, by affirming *p*, constitutively attempt to get it right whether *p*. Your belief is a constitutive attempt through means (affirmation) that would ground your success like Phil's stroke is, *mutatis mutandis*, when it comes to making the putt.

<sup>21</sup>Constitutive attempts can be assessed for success, competence, and aptness, where 'success' is indexed to whatever one (by taking constitutive means) constitutively attempts to secure.

make the basket, and Phil’s putt is a (constitutive) attempt to make the putt, we still haven’t said what makes a thinker’s *belief* a (constitutive) attempt to (broadly speaking for now) ‘get it right’ *vis-à-vis p*?

There are two ways to answer this question within telic virtue epistemology – what I’ll call the ‘easy way’ and the ‘hard way’. The ‘easy way’ involves no appeal whatsoever to any *intention* of the thinker; on this view, the kind of attempt one makes when one believes is related *teleologically* to correct (i.e., true) representation.<sup>22</sup> This way of thinking is very similar to the familiar idea (embraced by normativists about belief) that part of *what it is to believe* (rather than, say, to hope) is to take up a state that would be correct iff true.

The ‘easy story’ (which adverts to teleological rather than intentional aims) can in principle be supplemented with different substantive glosses. For example, one could give a further explanation here in terms of etiological functions.<sup>23</sup> Perhaps believing teleologically aims at truth because the etiological *function* of belief is accurate representation.<sup>24</sup> The important point, though, is that there is a ‘no frills’ version of telic virtue epistemology available which maintains that believers aim *implicitly* (though not intentionally) at truth (while maintaining neutrality on the substantive worth of this implicit aim). And then, as the thought goes, token constitutive attempts at truth (at which believers implicitly aim in affirming the relevant content) might be good (or not) with reference to evaluative norms of success, competence, and aptness.

Ernest Sosa’s own virtue epistemology was *initially* friendly to this easy story<sup>25</sup>, and it is this easy story that many epistemologists (e.g., John Greco) have in mind when they hold that knowledge is type identical with apt belief.<sup>26</sup> On this way of thinking: (i) Given belief’s teleological aim is truth, in believing, *we* (implicitly) teleologically aim at truth; and (ii) apt belief is a be-

<sup>22</sup>On the idea that functional states can have teleological aims, consider, e.g., Sosa’s remark that ‘a state of alertness in a crouching cat may be aimed at detecting vulnerable prey. Whether as a state it can count as a [sic. attempt] in any ordinary sense is hence irrelevant to our focus on [sic. attempts] that have an aim and to which we may then apply our AAA aim-involving normative account. All that really matters for this latter is that the entity have a constitutive aim, whatever may be its ontological status or the label appropriately applicable to it in ordinary parlance’ (Sosa 2015, Ch. 5, n. 5). For other discussions of functional and teleological assessment, see Sosa (2017, 71–72, 129–30, 152, 2021, 24–31, 52–58, 64, 110, 118).

<sup>23</sup>See, e.g., Graham (2014) and Simion (2019).

<sup>24</sup>This is but one optional further thesis one might embrace.

<sup>25</sup>In particular, I have in mind Sosa’s performance-theoretic epistemology *prior to Judgment and Agency* (2015). This includes, for example, Sosa (2007).

<sup>26</sup>Versions of this idea are found in Greco (2009, 2010), Turri (2011), Vargas (2016); and Carter (2016).

lief whose *correctness* (viz., *truth*) is creditable to, or because of, the thinker's competence.

*This* is how, on the easy story, the telic virtue epistemologist characterises the sense in which *epistemic normativity depends on our goals*. Epistemic normativity (at least, non-substantive, evaluative epistemic normativity) *depends* inextricably on our implicit, teleological aimings. In a bit more detail: evaluative norm facts – viz., facts about whether a thinker knows, justifiably believes, etc., – obtain in part in virtue of facts that include their aiming teleologically as they did. In slogan form: No knowledge and no justified beliefs without teleological aiming.

The problem is, this *can't* be the full story.

d. In order to see why the 'simple story' needs filled out further, just consider that at least sometimes – e.g., when we want to settle the matter of whether *p* – we are actually *intending* to do something more than *just get it right*.<sup>27</sup> We want to get it right *aptly*, through competence. *That* fact might explain why, for example, we spend all the time gathering evidence we do that bears on a particular question whether *p*, as opposed to (say) just make a hasty guess.

When the question is important or of special interest to us, we are trying to get it right *risk aversely*; this is very different from, by way of analogy, just throwing the ball at the hoop and hoping for the best. The basketball coach berates a player *even if they make a shot taken* if the shot incurs too much risk.<sup>28</sup> The basketball player is supposed to be *trying* to make it *not just any old way* (by means that incur any level of risk whatsoever), but to make it *aptly*, through competence that would minimise risk.

If that's right, then *in so far* as the normativity of *knowledge* is a special case of telic normativity (the normativity of attempts *as* attempts), we need some way of making sense of how knowledge-apt attempts (beliefs) might, *like basketball shots*, be attempts of the sort that aim at *aptness* and not merely at success understood as correctness. But we can't do that just yet, at least not on a picture on which believing always aims merely teleologically at *correct* representation.

e. Here is the final 'twist' in the tale: within the broad category of 'belief', we should really distinguish *two distinct kinds of (constitutive) attempts* we regularly make under this description. First, there is the kind of belief you have,

<sup>27</sup>Within virtue epistemology more generally, this idea was first developed in *Judgment and Agency* (Sosa 2015).

<sup>28</sup>See Sosa (2015, Ch. 3) and Carter (2016).

e.g., about your immediate surroundings, which guide your actions, but which aren't themselves answers to any 'whether  $p$ ' questions you've attempted to settle. For example, your belief (as you're driving a car) that you are approaching an intersection, or that a car has just passed you. These *functional* beliefs, important for successful practical reasoning and action, don't correspond with any question you might consciously consider and are thus not the answers to any such questions.

Compare these functional beliefs with *judgmental beliefs*, the sort that are the result of a deliberative inquiry into a question of whether  $p$ . When judging whether  $p$ , we aim – intentionally – to get it right through competence and not just through luck. On Sosa's more recent virtue epistemology – a version of which I embrace<sup>29</sup> – we can conceptualise a *judgment* as a constitutive (rather than instrumental) attempt to get it right *aptly* by affirming that  $p$ .

*Objection:* But wait! Doesn't this imply doxastic voluntarism of an objectionable sort? *Answer:* No. To say that judgment is a (constitutive) attempt, with *intention*, to attain a given aim (e.g., of apt belief) does not imply that how we judge is thereby under the sort of voluntary control whereby we could judge directly *through arbitrary choice*. In order to see why, let's consider what suffices for the kind of intentional action that judgment is. To do this, we can distinguish between a basic action and a 'simple' intentional action. Suppose I intentionally move a finger. Or think of a triangle. These are both deeds I do intentionally, but in each case, there's no other deed I do in the endeavour to do these things, either the physical basic action of moving a finger or the mental basic action of thinking of a triangle. Affirming that  $p$  is a basic action. When one affirms that  $p$ , one doesn't do this partly by doing something else in the endeavour to affirm that  $p$ . Basic action can be distinguished from 'simple' intentional action, where an agent aims to perform a deed (at  $t$ ) at least partly by performing a basic action,  $B$ , at  $t$ . Whereas affirming is a basic action, judgment is not; it is a simple intentional action.

The above sense in which judgment is a kind of (simple) intentional action, thus, doesn't imply at all that judging is voluntary in the sense that it could be reversed arbitrarily – viz., the sense in which researchers widely deny that belief, more generally, is voluntary.<sup>30</sup> Moreover, given that not all beliefs are judgmental, not all beliefs are intentional even in the above sense that does not imply voluntary control.

f. We've now distinguished between functional (teleologicalaly aimed) be-

<sup>29</sup>See Carter (2020) and, for a collective version of telic virtue epistemology, Carter (2020).

<sup>30</sup>See Williams (1970).

iefs and judgmental (intentionally aimed) beliefs.<sup>31</sup> Functional knowledge is teleologically aim dependent, judgmental knowledge is intentionally aim dependent.

This distinction allows us to now separate two kinds of theses at odds with epistemic instrumentalism:

**Weak telic instrumentalism:** Epistemic normativity (at least, non-substantive, evaluative epistemic normativity) *depends on our goals* in that it depends on our implicit, teleological aimings.

**Strong telic instrumentalism:** Epistemic normativity (at least, non-substantive, evaluative epistemic normativity) *depends on our goals* in that it depends on our conscious, intentional aimings.

Both weak telic instrumentalism and strong telic instrumentalism are incompatible with epistemic anti-instrumentalism, in that each thesis allows for epistemic normativity to depend on our goals *qua* aimings.

That said, those who go in for the *strong* thesis are perhaps more interestingly at odds with anti-instrumentalism.<sup>32</sup> Even though neither the strong nor the weak thesis permits substantive epistemic norms to depend on goals (in any sense), the strong thesis *but not the weak thesis* allows for at least some kind of epistemic norms to depend on *intentional* aims. This bears relevance because instrumentalists most inclined to draw the ire of the anti-instrumentalist (e.g., Foley, Kornblith), but *not* weak telic instrumentalists, allow for epistemic normativity to depend on intentional goals *qua* intentional aims.

g. In sum: *all* telic virtue epistemologists are minimally incompatible with epistemic anti-instrumentalism in virtue of a commitment to *at least* weak telic instrumentalism. This includes, e.g., Sosa (2007, Ch. 2; 2010, 471–71),

<sup>31</sup>This distinction between functional and judgmental belief first appeared in Sosa (2015), however, discussion of functional belief appears earlier, e.g., in Sosa (2007).

<sup>32</sup>While it is beyond what we can go in to in any detail here, it is worth noting that in the special case of judgment (of the sort that is countenanced by strong telic instrumentalists), there is a *further* normative dimension that comes into play when we theorise about what norms determine what *risks to aptness* a thinker in the market for fully apt judgment can non-negligently ignore. Different answers to this question are found in Sosa (2015, Ch. 3; 2020) and in the theory of fully apt judgment developed in Carter (2020). However, *despite* these differences, on both of these approaches to fully apt judgment, the fundamental explainers of epistemic status are telic norms.

Greco (2010, Ch. 5; 2003, 111), Turri (2011, 8–9) Zagzebski (1996, 271–72), Riggs (2002, 93–94), Carter (2016, 147–49), Littlejohn<sup>33</sup> (2014, 369), and Broncano-Berrocal (2017) However, there is space to move further away from anti-instrumentalism via strong telic instrumentalism (e.g., Sosa 2015, Ch. 3; 2021, Ch. 2; 2020, 3–4; Carter 2020, 32–34; 2020, secs. 3–5; Carter and Sosa 2021, sec. 2.1), and this is so *even though* neither telic instrumentalist thesis (neither weak, nor strong) commits one to the kind of instrumentalism (about *substantive* epistemic norms) that stands centrally in the anti-instrumentalist’s line of fire.

### 3. Zetetic norms and aim dependence

a. At this point it will be useful to return to the idea, central to both weak and strong telic instrumentalism, that epistemology focuses on objectives neutral in substantive normativity. As the datum of trivial truths suggests, truth is not *per se* substantively valuable – not any more than hitting a target with an arrow is *per se* substantively valuable – and neither is our knowledge of trivial truths.

But – as a potential objection might go – aren’t certain habits of *inquiry* substantively valuable (and disvaluable)? And if so, might there be at least *some* substantive epistemic norms over and above the non-substantive, telic norms (viz., success, competence, aptness) that telic virtue epistemology countenances?

b. Here it is helpful to draw a line between two entirely different kinds of norms that plausibly govern (broadly speaking) the moves we make as inquirers. On the kind of view I’ve been outlining here, the normativity of *knowledge* – of particular importance in epistemology – is a special case of (non-substantive) telic normativity. *Even so*, there seems to be another *kind* of normativity that bears on our intellectual lives. That is, there are plausibly substantive *zetetic* norms of intellectual ethics, norms that govern, e.g., what are *worthy* questions to take up in the first place, and how deeply to invest our finite resources in investigating them.<sup>34</sup>

Zetetic norms have an important *ethical* component (see Carter and Sosa 2021). For example, part of what explains the badness of the inquirer who

<sup>33</sup>It’s worth noting that Littlejohn’s (2014) defence of this view in this particular paper is centrally aimed at defending the proposal against fake-barn style counterexamples.

<sup>34</sup>For some different takes on zetetic norms and their place in epistemology, see, e.g., Friedman (2020, *forthcoming*) and Kelp (2021).

spends her time gobbling up trivial knowledge (suppose, she believes ‘*p*’ and then generates and affirms long disjunctions of the form  $\langle p \text{ or } q \text{ or } r \dots \rangle$ ) is that her inquiries are not *worthy* ones. Why not? This surely has something to do with substantive moral considerations, including what inquiries will contribute to a good life on the whole, perhaps also to the betterment of others.

Note, though, that substantive zetetic norms regulate not *only* what inquiries are worthy of pursuit, but also, for those inquiries we do choose to pursue, how *deeply* to pursue them – viz., at what point to terminate inquiry through affirmation one way or another.<sup>35</sup> To bring out this latter point, consider the recent literature on *moral deference*, where a common position holds that there is something problematic about terminating an inquiry into a moral proposition purely on the basis of another’s say so.<sup>36</sup> For instance, there is something amiss with believing that cruelty is wrong *just because your teacher told you so*, but for no other reason, and then giving no further consideration to the question. Such questions are worthy, morally speaking, of further investigation – we plausibly have (substantive) moral (normative) reasons to not terminate our inquiries into such questions on mere say so, even if mere say so (e.g., from a reliable source) is enough to get us testimonial knowledge.

c. So it is open, then, to the telic virtue epistemologist to accept that substantive normativity governs our inquiries – at least in the ways noted above. What is the significance of this concession? Is it somehow ‘going back’ on the core thesis that epistemic normativity is not substantive normativity?

No, but this is simply because the normativity of knowledge is different from the normativity of *intellectual ethics*, which is crucially important for regulating which inquiries to take up and how thoroughly to sustain them, but which is *irrelevant entirely* to our constitutive attempts at knowing. To emphasise this point, remember again the assassin’s shot might be an *excellent* shot (*qua* shot) and at the same time a moral abomination. That it is a moral abomination doesn’t undermine the telic credit to the assassin. Likewise, poor *intellectual* shot selection – by the lights of intellectual ethics – doesn’t undermine the telic credit to the thinker who *does* come to know a pointless truth, or who *does* come to know an important truth, worthy of deeper investigation, just on testimony.<sup>37</sup>

<sup>35</sup>For discussion, see Sosa (2021, Ch. 1) and Carter (2021).

<sup>36</sup>See, e.g., Jones (1999); Nickel (2001); Driver (2006); Hopkins (2007).

<sup>37</sup>That is, testimonial knowledge is no less *knowledge* in a case where (e.g., when the content of the knowledge is a moral proposition) by the norms of intellectual ethics, one

*Question:* are the substantive, moral norms of intellectual ethics – those that govern what inquiries are worthy of taking up and investing time in – *goal dependent*? The kind of virtue epistemology I’ve gone in for here rightly leaves this wide open, a question that requires ethical investigation, including about what counts as valuable ways to spend one’s finite time.<sup>38</sup> A stance here (even if important!) on this question in intellectual ethics is not going to be implied by our commitment to thinking that epistemic normativity – the normativity of knowledge – is telic normativity.

It is a metaepistemological question how exactly we should locate morally-infused zetetic norms within the *subject matter* of epistemology (see [Carter and Sosa 2021, sec. 2.1](#)). Some virtue epistemologists working in the responsibilist tradition are inclined to place zetetic norms at the heart of epistemology’s subject matter.<sup>39</sup>

For the present purposes, it suffices to register that telic norms (e.g., success, competence and aptness) are epistemic norms that are goal-dependent (in the weak and strong senses outlined), and this is so *regardless* of whether zetetic norms (wherever we should locate these in the subject matter of epistemology) are goal dependent – a point on which telic virtue epistemology *as such* is noncommittal.

#### 4. Where the dispute lies (and where it does not)

a. Let’s take stock. Here are some points of *agreement* between myself and Worsnip:

---

ought to have not terminated one’s inquiry on another’s say so. Of course, *were* one to have aimed at knowledge-first hand, or understanding (which might well be a matter of good intellectual ethics), and one then attained *that* aim aptly, one might then have a different kind of achievement, corresponding with this more demanding aim ([Sosa 2021, Ch. 1](#); [Carter 2021](#)). The crucial point here is that, *even so*, in a case where one (suppose, contrary to good intellectual ethics) *does* terminate a moral inquiry on another’s say so, their doing so isn’t precluded from being knowledge simply on account of the bad intellectual ethics, provided one’s testimonial belief in this case does well by the telic norm of aptness.

<sup>38</sup>By turning our attention to intellectual *ethics*, we not only consider what substantive norms might govern how to apportion our intellectual resources given our limitations, but also substantive norms that plausibly *prohibit* certain kinds of intellectual negligence. The moral deference discussion noted previously illustrates one example of this: there is plausibly a kind of negligence in terminating one’s inquiry into a moral question simply on someone’s say so. There are many other such examples however. To note one salient such example, consider examples of knowledge resistance ([Simion 2021](#)), especially in cases where one culpably avoids evidence gathering.

<sup>39</sup>See, e.g., Baehr ([2011](#)), Zagzebski ([1996](#)), Battaly ([2015](#)), Kidd ([2016](#)).

- Neither Worsnip nor I think that there are goal-dependent, substantive epistemic norms that govern what we ought to believe, or what is an (epistemic, normative) reason to believe what. In this respect, we both *disagree* with instrumentalists like Foley and Kornblith.<sup>40</sup>
- Both Worsnip and I take it that the kind of epistemic norms that are of central interest in epistemology are, in a certain sense, ‘sealed’ off from practical considerations. On my favoured approach, this is because the satisfaction of telic norms that govern our constitutive attempts at truth and knowledge aren’t in any way affected by practical considerations. For Worsnip, this is because only evidential considerations (and not practical goals) bear on what one ought to believe, and on what one has a reason to believe.<sup>41</sup>
- Both Worsnip and I agree that there are substantive norms that bear on how we conduct our intellectual lives, *and* that these substantive norms are plausibly not goal dependent. Worsnip construes the substantive norms that bear on how we conduct our intellectual lives as substantive *epistemic norms* on what we ought to believe (and what is a (normative) reason for what), whereas the approach I favour construes the substantive norms that govern our intellectual lives as *zetetic* norms (see §3), which regulate good and bad *intellectual ethics*, not our constitutive attempts at truth and knowledge.<sup>42</sup>

These are all important points of agreement – even if agreements reached on the basis of different reasons. Now, let’s transition into the murkier territory.

The kind of anti-instrumentalist position Worsnip embraces is committed to the view that epistemic normativity (and, again, here he has in mind substantive norms on what one ought to believe, and what one has a reason to believe) does *not* depend on our goals *qua intentional* aims. Here, the perspective of telic virtue epistemology from which I’ve been arguing is *divided within itself*. What I’ve called weak telic instrumentalism agrees with the anti-instrumentalist on this point – but *strong* telic instrumentalism is at odds with

<sup>40</sup>See Worsnip (Forthcoming, fn. 1) for some additional examples of epistemic instrumentalists he rejects.

<sup>41</sup>See Worsnip (Forthcoming, sec. 5).

<sup>42</sup>As a point of clarification: telic virtue epistemology is not as such committed to much about the *status* of ethical norms, other than that such norms are *substantive* (and imply *substantive* (normative) reasons) in a way that epistemic norms are not. Thus, my inclination to reject that zetetic, substantive norms of intellectual ethics are goal dependent doesn’t follow straightforwardly from my commitments to telic virtue epistemology.

the anti-instrumentalist here. For reasons I've indicated in §2(d-e), I'm sympathetic with the strong telic instrumentalist on this point. I think strong telic instrumentalism (in short) should be embraced – in addition to weak telic instrumentalism – in order to make sense of how our 'intellectual shots' can aim constitutively not only at accuracy but at (as judgmental beliefs do) *aptness*.<sup>43</sup>

b. The *key* point of disagreement – for the purpose of this exchange – concerns the *truth of anti-instrumentalism itself*. Both weak and strong telic instrumentalism (both theses of which I accept) are incompatible with anti-instrumentalism. Remember why this is so: if 'epistemic normativity', as it features in the anti-instrumentalist's core thesis, is interpreted widely enough to include *not only* substantive, prescriptive norms (such as norms about what we ought to believe on balance, and what is a reason to believe what), but also *non-substantive, evaluative norms*; and 'our goals' include telic aims, then (iii) from the core ideas at the heart of telic virtue epistemology, it follows that at least some epistemic norms depend on our goals, and *a fortiori*, epistemic anti-instrumentalism is false.

c. What if the anti-instrumentalist reasoned as follows: "Sure, epistemic anti-instrumentalism is false *if the thesis is interpreted so broadly* as to be a thesis about *all* epistemic normativity. However – contra the telic virtue epistemologist – the most important kind of epistemic normativity concerns what we ought (substantively) to believe, and what is a reason for what. And *these* norms do not depend on our epistemic goals. And so epistemic anti-instrumentalism is true in a restricted sense – viz., when restricted to the norms of particular importance in epistemology."

This envisioned reply, I think, helps to reveal what may be a further and arguably deeper metaepistemological divide between telic virtue epistemology (both in its weak and strong instrumentalist flavours) and the epistemic anti-instrumentalist – which is an *entirely different* divide, centred on different philosophical considerations, than what divides Worsnip and those instrumentalists who think substantive epistemic norms are goal dependent.

Whereas *those* instrumentalists will perhaps agree with Worsnip that our central attention in epistemology, when it comes to epistemic normativity, should be on substantive norms about what one ought to believe (and what is a normative reason for what), the telic virtue epistemologist resists this

<sup>43</sup>My preferred articulation of fully apt belief is defended in Carter (2020).

metaepistemological picture<sup>44</sup> – in line with the underlying thesis that unlike ethics, epistemology focuses on objectives neutral in substantive normativity.

Keeping this point in perspective lets us see how both the points of agreement *and* disagreement between weak and strong telic instrumentalists and anti-instrumentalists about the goal dependence of epistemic normativity are difficult to tease apart from what is perhaps a more fundamental, upstream dispute about the character of the kinds of norms that are of central importance in epistemology.

## References

- 10 Baehr, Jason. 2011. *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. OUP Oxford.
- Battaly, Heather. 2015. *Virtue*. John Wiley & Sons.
- Broncano-Berrocal, Fernando. 2017. “A Robust Enough Virtue Epistemology.” *Synthese* 194 (6): 2147–74.
- Carter, J. Adam. 2016. “Robust Virtue Epistemology as Anti-Luck Epistemology: A New Solution.” *Pacific Philosophical Quarterly* 97 (1): 140–55.
- . 2016. “Sosa on Knowledge, Judgment and Guessing.” *Synthese*, August.
- . 2020. “Collective (Telic) Virtue Epistemology.” In *Social Virtue Epistemology*, edited by Mark Alfano, Jeroen de Ridder, and Colin Klein. London: Routledge.
- . 2020. “De Minimis Normativism: A New Theory of Full Aptness.” *The Philosophical Quarterly*.
- . 2021. “Collateral Conflicts and Epistemic Norms.” In *Epistemic Dilemmas: New Arguments, New Angles*, edited by Kevin McCain, Scott Stapleford, and Matthias Steup. Routledge. <https://philarchive.org/rec/CARCCA-21>.
- Carter, J. Adam, and Ernest Sosa. 2021. “Metaepistemology.” In *Stanford Encyclopedia of Philosophy*.

<sup>44</sup>Note that this is not to say that for the telic virtue epistemologist *reasons* and *rationality* have no important place. For discussion on this point, see Sylvan and Sosa (2018).

- Driver, Julia. 2006. "Autonomy and the Asymmetry Problem for Moral Expertise." *Philosophical Studies* 128 (3): 619–44.
- Foley, Richard. 1987. *The Theory of Epistemic Rationality*. Harvard University Press.
- Friedman, Jane. forthcoming. "Zetetic Epistemology." In *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere*. Routledge.
- . 2020. "The Epistemic and the Zetetic." *Philosophical Review* 129 (4): 501–36.
- Geach, Peter T. 1956. "Good and Evil." *Analysis* 17 (2): 33–42.
- Glüer, Kathrin, and Åsa Wikforss. 2009. "Against Content Normativity." *Mind* 118 (469): 31–70. <https://doi.org/10.1093/mind/fzn154>.
- Graham, Peter J. 2014. "Functions, Warrant, History."
- Greco, John. 2003. "Knowledge as Credit for True Belief." In *Intellectual Virtue: Perspectives From Ethics and Epistemology*, edited by Michael DePaul and Linda Zagzebski, 111–34. Clarendon Press.
- . 2009. "Knowledge and Success from Ability." *Philosophical Studies* 142 (1): 17–26.
- . 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge University Press.
- Hopkins, Robert. 2007. "What Is Wrong with Moral Testimony?" *Philosophy and Phenomenological Research* 74 (3): 611–34.
- Jones, Karen. 1999. "Second-Hand Moral Knowledge." *The Journal of Philosophy* 96 (2): 55–78.
- Kelp, Christoph. 2021. *Inquiry, Knowledge, and Understanding*. Oxford: Oxford University Press.
- Kidd, Ian James. 2016. "Charging Others with Epistemic Vice." *The Monist* 99 (2): 181–97.
- Kornblith, Hilary. 1993. "Epistemic Normativity." *Synthese* 94 (3): 357–76.
- Kvanvig, Jonathan. 2008. "Pointless Truth." *Midwest Studies in Philosophy* 32 (1): 199–212.

- Littlejohn, Clayton. 2014. “Fake Barns and False Dilemmas.”
- McHugh, Conor. 2012. “The Truth Norm of Belief.” *Pacific Philosophical Quarterly* 93 (1): 8–30.
- Nickel, Philip. 2001. “Moral Testimony and Its Authority.” *Ethical Theory and Moral Practice* 4 (3): 253–66.
- Riggs, Wayne D. 2002. “Reliability and the Value of Knowledge.” *Philosophy and Phenomenological Research* 64 (1): 79–96.
- Shah, Nishi. 2003. “How Truth Governs Belief.” *The Philosophical Review* 112 (4): 447–82.
- Shah, Nishi, and J. David Velleman. 2005. “Doxastic Deliberation.” *Philosophical Review* 114 (4): 497–534. <https://doi.org/10.1215/00318108-114-4-497>.
- Sharadin, Nathaniel. 2018. “Epistemic Instrumentalism and the Reason to Believe in Accord with the Evidence.” *Synthese* 195 (9): 3791–3809.
- Simion, Mona. 2019. “Epistemic Norms and Epistemic Functions.” *Book Manuscript*.
- . 2021. “Resistance to Evidence and the Duty to Believe.” *Philosophy and Phenomenological Research*.
- Simion, Mona, Christoph Kelp, and Harmen Ghijsen. 2016. “Norms of Belief.” *Philosophical Issues* 26 (1): 374–92.
- Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford University Press.
- . 2010. “How Competence Matters in Epistemology.” *Philosophical Perspectives* 24 (1): 465–75.
- . 2015. *Judgment & Agency*. Oxford University Press UK.
- . 2017. *Epistemology*. Princeton: Princeton University Press.
- . 2020. “Default Assumptions and Pure Thought.” *Manuscript*.
- . 2021. *Epistemic Explanations: A Theory of Telic Normativity, and What It Explains*. Oxford University Press.
- Sylvan, Kurt, and Ernest Sosa. 2018. “The Place of Reasons in Epistemology.” In *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star.

- Turri, John. 2011. "Manifest Failure: The Gettier Problem Solved." *Philosopher's Imprint* 11 (8). <http://hdl.handle.net/2027/spo.3521354.0011.008>.
- Vargas, Miguel Ángel Fernández. 2016. *Performance Epistemology: Foundations and Applications*. Oxford University Press.
- Wedgwood, Ralph. 2002. "The Aim of Belief." *Philosophical Perspectives* 16: 267–97. <https://doi.org/10.1111/1468-0068.36.s16.10>.
- Williams, Bernard. 1970. "Deciding to Believe." In *Problems of the Self*, edited by Bernard Williams, 136–51. Cambridge University Press.
- Worsnip, Alex. Forthcoming. "Epistemic Normativity Is Independent of Our Goals." In *Contemporary Debates in Epistemology*, edited by Blake Roeber, Matthias Steup, John Turri, and Ernest Sosa, 3rd ed. Wiley-Blackwell.
- Zagzebski, Linda T. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press.