

Objectivity and Triviality

Billy Dunaway
University of Oxford

Draft of 10th October 2014

1 A question and its methodology

Is objectivity trivial? In the present section I set out the question in more detail, and lay down some methodological ground rules for addressing it. Then in §§2-4 I make three related points that bear on whether an affirmative answer to the question should be given.

'Objectivity' is a term of art, and so pinning down what exactly is at issue in debates over objectivity may be a fruitless endeavor. One may choose to use a term of art however one wishes. So an answer to the question of what is at issue ought not to take the form of arbitrary legislation requiring that the term be used in certain ways. Fortunately many philosophers will have a decent working grasp of the term 'objective' and will agree in its application to a wide range of cases (I discuss some apparently clear-cut cases below). We could then attempt to ask what properties all of these cases have in common, and conclude that 'objective' picks out one of these properties. However there are several problems with this approach by cases. (i) Inconsistent intuitions: firmly held intuitions on the matter may be inconsistent with one another—some philosophers may find it natural to say that a particular view denies objectivity to a domain, while others disagree. (ii) Gerrymandered intuitions: Intuitions concerning objectivity may be sensitive to multiple highly natural kinds. (iii) No new predictions for new cases: if we allow highly gerrymandered properties as candidate referents for 'objective', then when confronting a case on which intuitions deliver no clear verdict, it will be unclear how to settle the matter. There will be multiple gerrymandered properties that fit all of the clear-cut cases, and some but not all of these will be instantiated by the new case. This is just a sketch, but it suggests several reasons why simply consulting intuitions about cases is not a promising methodology when investigating a term of art.

A more promising approach, which I will follow here, is to ask what structural features are suggested by common uses of 'objective', and whether there is any relatively natural and ungerrymandered kind that has these structural features. Here cases can be of some help: with a healthy list of examples before us, we can formulate some general roles that the notion is supposed to play. These structural features can then serve as the starting point of a theoretically respectable inquiry

into the target notion: we can ask whether there is any relatively natural kind that plays these roles. Any property that satisfies these criteria will then have a good claim to being at least *one* legitimate notion of objectivity.

There are several possible outcomes to this type of investigation. First there might be multiple sets of plausible structural features, and a fairly natural kind that satisfies each set. Each of these properties will then have a claim to being *a* legitimate notion of objectivity: there is no point in insisting that one of the candidate notions is the ‘right’ meaning for a term of art, as nothing is gained by engaging in linguistic stipulation of this kind. Second it is an open possibility that one or more of the relevant theoretical roles are not satisfied by any very natural property. Our practices with a term of art might not be tracking any theoretically interesting kind. In this case, we should conclude that where there seemed to be a theoretically interesting property, none in fact existed.

As a starting point for this type of investigation, some cases will be helpful. Ethics is one subfield in philosophy that provides a rich catalogue of intuitive examples of objective and non-objective views. One example of a view on which ethics is not objective is a Subjectivism, a simple version of which holds that a statement of the form ‘*a* ought to ϕ ’ is true just in case *a* approves of ϕ -ing.¹ Another example is a Constructivist view, which to a rough first approximation holds that ‘*a* ought to ϕ ’ is true just in case were *a* to reflect on all of her practical commitments, she would judge that she ought to ϕ (I will discuss refinements of the Constructivist idea in later sections). These views are to be contrasted with a non-naturalist view on which obligation is an unreducible, *sui generis* component of reality, or naturalist view like that of Railton (1986), on which obligation can be reduced to shared psychological features.

There are examples from other domains as well. While ethics is a domain for which it is an open question whether the right view is an objective one or not, there are other domains for which it is commonly felt that the only live options are non-objective views. Examples include the comic, taste, and the domain of legality. What exactly it is that makes these domains non-objective is a question yet to be answered. But whatever separates these domains off as non-objective, it must not be a property that paradigmatically objective domains such as physics or biology also have.

Since I am treating it as a substantive question whether a particular view about any of these domains makes them out to be objective, it isn’t a priori that the intuitive judgments about the examples above will turn out, on investigation, to be correct. It may be that no very natural property both applies to the cases and satisfies any of the distinctive structural features of objectivity. In particular, I won’t assume that it is true *by definition* that the Subjectivist or Constructivist views entail ethics to be non-objective. Of course it would be trivial to find a highly unnatural and gerrymandered referent for ‘objective’ that licensed our intuitive

¹An ancestor of the Subjectivist view is the “Emotivism” of Stevenson (1937); see also Harman (1975) for a more sophisticated version of the view.

judgments about these cases: we could talk about the property named by the predicate 'is either accepted by theorists who are labelled 'Constructivists' or is accepted by theorists who are labelled 'Subjectivists' ...'. But this would be an uninteresting property, underserving of the debates that objectivity has inspired. Thus we may not be able, after investigation, be able to vindicate the judgment that Constructivism and Subjectivism are non-objective views. The triviality of objectivity remains a live option.

I will address three purported structural features of objectivity in the rest of this paper. The first (§2) revolves around the notion of a "shortcoming-free disagreement". According to some, what is distinctive of non-objective domains is that they support the possibility of disagreements over their subject-matter where neither of the disputants is "at fault" in some sense. The second (§3) is what I call a "collapse" principle, according to which, for any candidate characteristic of non-objective domains Φ , if it is objective that a domain has Φ , then the domain in question is wholly objective. The third (§4) is the support of "easy knowledge". That is, non-objective domains are supposed to make knowledge of that domain easy in a distinctive sense, avoiding epistemological puzzles that plague robustly realist, "objective" views of the same domain.

Before proceeding to these questions, a few additional structuring assumptions about objectivity are in order.

We said earlier that 'objective' may be used with a number of different senses. It will be helpful here to circumscribe the candidate senses that we will be interested in. This is not because the other senses are uninteresting, but only because an investigation of this kind is more effective if it treats one sense at a time. First: various authors have taken the bearers of objectivity to be very different kinds of things. The official position in Nagel (1989), for instance, is that objectivity is primarily a property of a methodology—ethics is objective on this proposal if the method of ethical investigation takes there to be a further question about whether a set of practical commitments is correct or not.² Second: there is also disagreement over what kind of property objectivity is. For instance, there are uses on which 'objective' is intended to pick out a metaphysical property, perhaps having to do with mind-independence in some sense. And there are others on which it is intended to pick out a semantic property, to be contrasted

²Here is Nagel:

Objectivity is a method of understanding [...] To acquire a more objective understanding of some aspect of life or the world, we step back from our initial view of it and form a new conception which has that view and its relation to the world as object. In other words, we place ourselves in the world that is to be understood. The old view then comes to be regarded as an appearance, more subjective than the new, and correctable or confirmable by reference to it. (Nagel 1989, 4)

with properties for whom a relativist semantics is correct.³

In what follows treat ‘objective’ as if it picks out a property of propositions. And I will be interested in objectivity as an epistemic kind, i.e., a property a proposition has in virtue of its connections to knowledge, belief, and other epistemic notions. In doing so, I will be primarily engaging with a tradition of thought according to which views that deny the objectivity of ethics thereby gain distinctive advantages in epistemology. For instance, a realist view on which ethics is fully objective has been accused in various ways of making it difficult to explain how we could ever come to know that objective facts in question. (The absence of any causal connections between the normative and our belief-forming processes, on many versions of this argument, is supposed to indicate the suspect epistemic status of these facts.) And the Constructivist and Subjectivist views outlined above are commonly thought to avoid this problem. Whatever their advantages or disadvantages in other areas, these views are supposed to have a distinctive advantage with respect to the epistemology of ethics. This thought will be developed in more detail below; here I only give a rough sketch of the epistemic property (to be contrasted with metaphysical and semantic properties) we will be interested in. I will call this the *epistemic conception* of objectivity, and I will ask in what follows what structural features a fairly natural epistemic kind that goes under the heading of ‘objectivity’ might plausibly be expected to satisfy.

The rest of this paper addresses a set of related questions which all revolve around the potential triviality of objectivity. The thesis that objectivity is trivial I will take to be the thesis that it follows from minimal uncontroversial assumptions that, for any domain D whatsoever, the facts in D are objective. Some of the structural features I investigate below (§3) quite clearly entail that objectivity is trivial *if* they are indeed structural features of the notion. However, I will argue that there are strong reasons for rejecting that the feature in question really does characterize objectivity. Others (especially those discussed in §2 and §4) less obviously entail that objectivity is trivial if they are structural features of the notion. But I will outline some arguments for the thesis that they do have this consequence. I do not take a final stand on the truth of this thesis that these in fact are structural features of objectivity, but am rather concerned with what follows concerning the triviality of the notion *if* they do in fact characterize it.

2 Shortcoming-free disagreement as a structural feature

A seminal work on the epistemic conception of objectivity can be found in Wright (1992). Wright develops the intuitive idea that non-objective facts are those for which there is a possibility of a “shortcoming-free disagreement”: when dealing with objective facts, subjects who form disagreeing opinions are such that one of them must display some kind of cognitive shortcoming. The paradigm case is

³See for instance discussions in Rosen (1994) and Cappelen & Hawthorne (2009) for discussions revolving around metaphysical and semantic notions in the area.

perception of ordinary objects; when two subjects disagree about whether there is a chair in the room in which they are both present, it follows that one of them is epistemically non-ideal, by virtue of having impoverished visual inputs, improperly functioning perceptual apparatus, bias, etc. Non-objective facts, by contrast, allow for the possibility of shortcoming-free disagreement: it is in the nature of these facts that some perfectly competent agents with no epistemic shortcomings may nonetheless disagree.

Wright's motivating case for this rough sketch is the comic: judgments about what is funny are supposed to be paradigmatic examples of judgments about a domain that is non-objective. According to the rough theory on offer here, the non-objectivity of the comic shows up in our epistemic lives in the form of the possibility of shortcoming-free disagreement. Two cognitive agents who arrive at distinct judgments concerning whether a particular joke is funny, need not be such that one is not fully competent, or not in epistemically friendly environments, etc. Instead, it is compatible with their disagreement that they are *both* in the same epistemic circumstances by virtue of each agent meeting the same standards for ideality. To insist otherwise, by claiming that one of the disagreeing agents must exhibit some shortcoming, would be to chauvinistically project a specific comic sense on those who do not share it.⁴

The terminology used to formulate this intuitive idea is at best suggestive. Wright recognizes this, and develops at length a proposal that improves on our use of 'epistemically ideal' and the like. This is critically important for developing a satisfactory proposal: any theoretically interesting kind that is a candidate referent for a term of art will need to bear substantive and workable connections to other important concepts. It would be little use to conduct an investigation into whether ethics is objective if the question were not formulated in a way that permitted a determinate answer or left the status of epistemically important notions like knowledge unclear. Wright takes some steps in this direction by offering the notion of *Cognitive Command* as a precisification of this rough but intuitive characterization of objectivity:

A discourse exhibits Cognitive Command if and only if it is a priori that differences of opinion arising within it can be satisfactorily explained only in terms of "divergent input", that is, the disputants' working on the basis of different information (and hence guilty of

⁴There are substantial questions about how to implement the characterization of neither judge being wrong in such a case. Dummett (1963) famously held with respect to the related issue of realism that, in domains of which irrealism holds, the law of bivalence fails. But few would find the idea that the law of non-contradiction fails for non-objective domains to be plausible. Though a naive reading of the 'neither party has made a mistake' locution might suggest that both parties must hold *true* beliefs, I will assume that this consequence can be avoided. The mechanism by which this is to be done is not something I will explore in detail here, but I assume that this can be done by appeal to some relativist or contextualist semantics for judgments about the comic, but will not press the point here. See Cappelen & Hawthorne (2009) and Kölbel (2002) for more discussion of these issues.

ignorance or error, depending on the status of that information), or “unsuitable conditions” (resulting in inattention or distraction and so in inferential error, or oversight of data and so on), or “malfunction” (for example, prejudicial assessment of the data, upward or downwards, or dogma, or failings in other categories already listed). (Wright 1992, 93)

The hypothesis on offer, then, is that objective domains are those which are such that discourse about them exhibits Cognitive Command.⁵

2.1 *Is Cognitive Command trivially satisfied?*

There are in general difficulties in making sense out of the notion of a shortcoming-free disagreement, and (as I will argue below) Wright’s sharpening in terms of cognitive command is no exception.⁶ The literature has focused on explaining how the notion of shortcoming-free disagreement is not incoherent in virtue of the role *truth* plays in the characterization of the notions of an epistemic shortcoming and disagreement.⁷ In short the worry is this: one might worry that

⁵There is an argument in Dworkin (1996) that ethics will be trivially non-objective on Wright’s criterion. Dworkin interprets Wright’s cognitive command criterion as equivalent to the following first-order normative claim:

[P]eople have a reason for asserting that there must be such an explanation of a particular kind even when they have no idea what it is. (Dworkin 1996, 106)

That is, whenever a domain is objective, we will have a reason to assert that there is a cognitive shortcoming but not have any evidence as to shortcoming in particular is present. But, he goes on to say, it is false that we have such reasons:

People have no reason to claim that those who disagree with them must lack some information they have, or suffer from some intellectual incapacity or character defect, when they have no evidence of any such ignorance or incapacity or defect. That claim, in those circumstances, would be empty rhetoric . . . (Dworkin 1996, 106-7)

Thus the claim that ethics is objective is a false normative claim.

However it is highly questionable that, as Dworkin assumes, one needs to have evidence about that a particular defect of some kind is present in order to have a reason to assert (or even believe) that *some* defect is present. Suppose God tells you that he will flip an objectively random coin and, if it lands heads, will give Sue misleading perceptual inputs and, if it lands tails, will rewire her visual system and thereby cause it to process inputs in a highly unreliable manner. Before knowing how the coin lands, it seems that one has a great reason to assert (or believe) that Sue has exhibited some shortcoming, but that one has no reason to assert that the shortcoming is in particular misleading inputs or faulty processing.

⁶The possibility of shortcoming-free disagreements is also at the center of the recent epistemological literature on “peer disagreement” as by definition ‘epistemic peers’ are supposed to be alike in all epistemic respects but nonetheless capable of disagreeing in some circumstances. See Hawthorne & Srinivasan (2013) for a similar point with respect to peer disagreement.

⁷Wright summarizes this line of argument as follows:

...it is a priori that any difference of opinion concerning the comic, when not attributable to vagueness and so on, must involve cognitive shortcoming, since, if all else fails, ignorance or error will at least be involved *concerning the truth value of the disputed statement*. (Wright 1992, 149)

any dispute will trivially involve no shortcoming-free disagreement since (i) if two disputants disagree, then they hold beliefs which cannot simultaneously be true, but (ii) if one of the disputants holds a false belief, then she thereby manifests an epistemic shortcoming.

Thus a clarification of Cognitive Command that serves Wright's purposes rests in part on whether there is a workable notion of what a cognitive shortcoming is which does not follow from merely believing falsely.⁸ However I will bypass this debate and assume that the Cognitive Command condition can be explicated in such a way that does not entail that the mere fact that one party has a true in any dispute entails that a shortcoming is present. Still, I will argue, the component conditions of Cognitive Command are trivially satisfied so long as we make the additional anti-skeptical assumption that one can know paradigmatically non-objective claims.

I will consider each of the component conditions of Cognitive Command in order.

Divergent input: Begin with the notion of "divergent input". For simplicity, let's assume for the moment a dispute for which the other component conditions are not satisfied, and so neither party in the dispute in question is in unsuitable conditions, or has displayed a cognitive malfunctioning (we will revisit this assumption below). It is plausible then to gloss the divergent input condition as one which obtains when the two parties to a dispute are working on the basis of different information, which is to say they have different evidence.

But on one very natural sharpening of what one's evidence is, one's evidence is the total body of what one knows. (*cf.* Williamson (2000, Ch. 9)) Thus given this sharpening, a dispute that does not involve divergent input will have to be one in which both of the disputants have exactly the same knowledge. But it will be trivial that (given our other assumptions) any case of disagreement will not be such a case, since exactly one of the disputants will have knowledge in the disputed area. Suppose there is a dispute over p ; exactly one disputant has a true belief p . Moreover if they have formed the belief without cognitive malfunctioning and in friendly epistemic circumstances, it will be very natural to say that they *know* that p . But since the other disputant disagrees on the matter, it follows that this is a case of divergent input. The other party does not believe p , as they are the other party in a dispute over the truth of p . Hence they do not know p , and the inputs diverge. And the conclusion is perfectly general, as we needed no assumptions about which domain p belongs to, only that the facts at issue are knowable and are in fact known so long as the claim believed is true and the other

See also Sainsbury (1996), Kölbel (2004), and Wright (2001).

⁸A semantic relativist in the style of Kölbel (2004) and MacFarlane (2007) will the charge of incoherence by rejecting (i): in a disagreement, the disputants might each believe something which is *true from her own perspective*. Believing what is true from one's own perspective need not involved any epistemic shortcoming. Semantic relativism is not Wright's preferred solution to the issue. See also Cappelen & Hawthorne (2009, 129 ff.) for discussion of "faultless disagreement" cases as motivation for relativistic semantics.

conditions for responsible belief-formation are met. This makes the “divergent input” condition trivial.

Malfunctioning: Of course some will wish to opt for a less stringent conception of what it is for the two disputants to have the same evidence, so allow (without opting for skepticism) that it is sensible to claim that both parties in a disagreement have the same evidence. Here is one option: the disputants are such that, prior to forming beliefs about the disputed proposition p , their respective bodies of evidence (i.e., what they know) each make it equally likely that p is true. That is, if a and b are the disputants, and E_a is what a knows prior to forming an opinion as to whether p , and E_b is what b knows prior to forming an opinion as to whether p , then $Pr(p | E_a) = Pr(p | E_b)$. This suggests one notion of a cognitive malfunction:

One exhibits a cognitive malfunction by believing p just in case one believes p and p is not very likely on the evidence one had prior to coming to believe p .

However while this approach salvages the divergent input condition, there is a very natural reading on which it entails that one of the disputants has malfunctioned. For (in order for the dispute to be genuine) at least one of disputants must have believes $\neg p$ when p is likely on her evidence. If we do not have a case of divergent input, then both disputants must have the same evidence, E . It cannot be that both p and $\neg p$ are very likely conditional on E . Thus in the above sense at least one of the disputants malfunctions in their belief-forming process.

One might reasonably resist a picture on which believing something that is unlikely on one’s prior evidence is sufficient for cognitive malfunctioning.⁹ But the complications will not help rescue a non-trivial reading of Cognitive Command. Suppose there are cases where believing what is improbable on one’s prior evidence isn’t a cognitive malfunction because one can know the proposition in question if one does come to believe it. This suggests an alternative notion of a cognitive malfunctioning:

One exhibits a cognitive malfunction by believing p just in case one believes p and one is not in a position to know p by coming to believe p .

But then, for the same reasons that made divergent input trivially satisfied on one natural understanding of “evidence”, one of the disputants must exhibit a cognitive malfunction in the above sense. The upshot here is that even if we

⁹In some cases one can come to know a proposition that is unlikely on one’s evidence—these are cases of “improbable knowing” as described in Williamson (2011). And there are mundane cases where, on one’s prior knowledge, it was unlikely that one would be in a circumstance c but, upon finding oneself in c , one can come to know p via perception. In both cases it is unlikely on one’s priori knowledge that, if one were to believe p , one would know it, but nevertheless one still acquires the relevant knowledge.

can find an acceptable sense of ‘evidence’ on which two disputants might share the same evidence, they will not fare equally well concerning the quality of their cognitive processes.

Unsuitable conditions: It is worth noting that there are other ways of making sense of the notion of a cognitive malfunction which render the unsuitable conditions component to Cognitive Command trivially satisfied. One might gloss malfunction in terms of a disposition: some belief-forming dispositions reliably (though not inevitably) manifest themselves by outputting beliefs that constitute knowledge; one might think of a belief as the result of a malfunction just in case it is not the result of such a disposition:

One exhibits a cognitive malfunction by believing p just in case one believes p and one’s belief is not the product of a belief-forming disposition that reliably produces knowledge.

It makes sense to say that, while two disputants’ beliefs were the product of the same reliably knowledge-producing disposition, in only one of these manifestations of this disposition, on the occasion of the dispute in question, produced knowledge. So it is coherent to say, of a single dispute, that the disputants had the same knowledge prior to forming an opinion in a dispute, and that their (divergent) opinions were formed as the product of the same generally reliable belief-forming disposition. But it follows from this that one of the disputants was in especially unfriendly epistemic circumstances. This is because it follows from this that one of the disputants (i) had all of the evidence as someone who came to know whether p , and (ii) formed a belief by manifesting the exact same disposition as the person who came to know whether p , yet (iii) formed a false belief as to whether p . It is very natural to say that someone who satisfies (i)-(iii) is in an especially unfriendly belief-forming environment.

Thus there is a plausible case that *if* knowledge is possible in a domain, Cognitive Command will be trivially satisfied. Crucially this relied on nailing down precise notions of the components of the Cognitive Command criterion. But with these precisifications in place, we can show that for any dispute as to whether p , if it is knowable whether p , then one of the disputants must have different evidence, exhibit a cognitive malfunction, or be in unsuitable epistemic circumstances.¹⁰

¹⁰Wright (2001, 60) considers an argument for a similar conclusion, which shows that in any dispute one agent must lack knowledge. His solution is to accept the argument, and to embrace skepticism for non-objective domains. He holds that it can never be affirmed that one party has exhibited a cognitive shortcoming because it can never be affirmed that any party has knowledge of facts in the domain that the other lacks. Disagreements about non-objective p , on Wright’s view, hold when the disputants do not know whether or not p , moreover also cannot say that they do not know since they also do not know that it is impossible to know whether or not p . (Wright 2001, 92)

Before moving on, some methodological points are in order.

Of course many theorists will find attractive precisifications of the crucial notions other than those I have considered here. Some may hope to use the notions of ‘justification’ or ‘rationality’, thinking they are well-suited to coherently formulate what is involved in a shortcoming-free disagreement: even if the disputants in a shortcoming-free disagreement cannot both have knowledge, they might both be equally justified in their beliefs. (Schafer (2011) suggests something like this view.) I will not argue here that this approach must be unsuccessful but instead wish to situate it in the methodological constraints outlined in §1. First, it needs to avoid treating the justification-centric ideology as primitive, unexplained theoretical terminology: since the present context is an investigation of the term of art ‘objective’, taking the justification-centric terminology to be primitive would merely swap one term of art for another. But in explaining the crucial notions it must also avoid an analysis of the ideology that makes it out to be about a highly gerrymandered, unnatural kind (otherwise the resulting analysis of objectivity would imply that the notion is not, in its epistemic sense, a very interesting one). I will not argue here that these *desiderata* cannot be met. I have considered some sharpenings of ‘evidence’, ‘malfunction’, and ‘unsuitable condition’, and on these sharpenings it is trivial that no dispute will satisfy all of the conditions for a shortcoming-free disagreement.

2.2 *A way forward*

Even those who are optimistic about a coherent notion of a shortcoming-free disagreement may nonetheless wish not to place too much weight on the notion in a discussion of objectivity. Although Wright opts for the shortcoming-free disagreement ideology in his characterization of Cognitive Command, his motivating picture behind an epistemic conception of objectivity does not make use of the *inter*-personal notion of a disagreement at all. Instead he says:

[I]n discourses of the relevant kind [i.e., about non-objective subject-matter], we are dealing with matters which essentially *cannot* outrun our appreciation: that there is no way in which something can be delicious, or disgusting, or funny, or obscene, etc., without being appreciable as such by an appropriately situated human subject because those matters are, in some very general way, constitutively dependent upon *us*. What we—most of us—find it natural to think is that disputes of inclination typically arise in cases where *were* there a ‘fact of the matter’, it would have to be possible—because of the constitutive dependence—for the protagonists to know of it. (Wright 2001, 59)

There is nothing in this motivating picture concerning disagreement, and more generally the interpersonal component to Cognitive Command is entirely lacking. This suggests that we can get at the core idea behind the epistemic conception

of objectivity with out getting bogged down in the issues that make a coherent formulation of a shortcoming-free disagreement.

At a first pass a plausible characterization of the epistemic conception of objectivity, as suggested by Wright, is the following: for non-objective domains, there is an especially tight connection between forming belief in hospitable circumstances and epistemic success. In other words, one isn't at risk of epistemic failure so long as one forms one's belief in certain conditions.

It is easy to see why one would be tempted to use the notion of a shortcoming-free disagreement to flesh this idea out: forming beliefs in the right circumstances guarantees epistemic success, then we might say of any two disagreeing subjects (in the right circumstances) that neither of them has exhibited any epistemic shortcoming. But while one prima facie plausible way of fleshing out the motivating idea behind an epistemic conception of objectivity focuses on the inter-personal notion of disagreement, this aspect isn't *essential* to the notion, as there is nothing in the concept of a connection between belief-formation in a domain and cognitive success in that domain that requires the inter-personal detour.

An extremely simple alternative characterization of an epistemic notion of objectivity in this mould, we can make use of the notion of a *belief-knowledge connection*. To avoid gerrymandering the proposal from the outset, we can take a simple view of when a subject has achieved epistemic success, and hold that the relevant success is just having *knowledge* of the claims in question. And as a starting point for a simple characterization of the conditions under which success is guaranteed, we can take the relevant conditions to be logical and probabilistic competence on the part of the belief-forming subject. (Clearly if relevant beliefs were arrived at via deductive incompetence or a failure to be probabilistically coherent, then they would not constitute knowledge even in paradigmatically non-objective domains.) According to the simple condition, then, it is necessary that any time a logically and probabilistically competent subject comes to believe that p , she thereby knows that p . That is, non-objectivity for a domain D requires the following:

BELIEF-KNOWLEDGE CONNECTION Necessarily, for any claim p in D and any subject s , if s is a logically and probabilistically competent subject, then if s comes to believe p , s thereby knows p .¹¹

A subject who is logically and probabilistically coherent is just one who deduces the logical consequences of what she believes, updates her credences by conditionalization and conforms to the laws of probability, and so on. **BELIEF-KNOWLEDGE CONNECTION** is, we cautioned above, an extremely simple gloss on the basic

¹¹One could try to argue that **BELIEF-KNOWLEDGE CONNECTION** is a plausible sharpening of the imagery involved in other framings of debates about objectivity. For instance, one might see it as a consequence of the metaphor of "imposing" concepts in Putnam (1987). Dummett (1978) articulates something like a connection of this kind, restricted to believers who have the relevant evidence, when he says:

idea and on its own will very likely lead to triviality. For example, logical and probabilistic coherence will not guarantee that on a Subjectivist view someone who forms an ethical belief will thereby come to have knowledge. For such a person may falsely believe that Jones in fact disapproves of murdering, and hence come to believe but not know that it is wrong for Jones to murder.

Thus while BELIEF-KNOWLEDGE CONNECTION has the virtue of being very simple, it is too simple to yield a non-trivial notion of objectivity. However it might provide a template for filling out similarly interesting notions that also have the virtue of being plausibly being non-trivial. In particular, we might supplement the relevant conditions to include more than deductive and probabilistic competence. We can obtain different (and potentially more substantive) versions of the BELIEF-KNOWLEDGE CONNECTION characterization of non-objectivity by expanding the conditions in different ways. Thus different epistemic conceptions of objectivity result from different approaches to filling in the following schema:

BELIEF-KNOWLEDGE CONNECTION SCHEMA (BKCS) Necessarily, for any claim p in D and any subject s , if s satisfies conditions C , then if s comes to believe p , s thereby knows p .

Of course BKCS is just a schema for a substantive characterization of non-objectivity, and is not itself such a characterization. There is still a possibility that there it is not possible to fill in the schema without gerrymandering the notion of objectivity to a degree that makes it a theoretically uninteresting property. I will have more to say about instances of the schema in the final sections (§§4-5). But first, the condition in schematic form can be used to assess a proposed structural feature of objectivity which I call the “collapse principle”. I will discuss this in the next section.

3 The collapse principle

Any discussion of the notion of objectivity will have to confront the central question in Gideon Rosen’s “Objectivity and Modern Idealism: What is the Question?”, which contains an important argument for the claim that objectivity is trivial. There are many interesting strands to Rosen’s discussion, and I cannot cover all of

Realism I characterise as the belief that statements of the disputed class possess an objective truth-value, independently of our means of knowing it: they are true or false in virtue of a reality existing independently of us. [...] The anti-realist insists on the contrary, that the meanings of these statements are tied directly to what we count as evidence for them, in such a way that a statement of the disputed class, if true at all, can be true only in virtue of something of which we could know and which we should count as evidence for its truth. (Dummett 1978, 146)

I won’t argue that BELIEF-KNOWLEDGE CONNECTION is in fact an acceptable sharpening of these views here; I only claim that *Belief-Knowledge Connection* articulates a sharpening of one of the plausible structural features of non-objectivity, and leave it as an open question whether it constitutes a satisfactory sharpening of Putnam’s and Dummett’s pictures as well, or whether these pictures can be made sufficiently precise at all.

them here. Instead I will only discuss one highly general “master argument” for this pessimistic conclusion. One instance of Rosen’s argument that objectivity is trivial can be found in his discussion of a “response-dependence” characterization of objectivity. On this characterization the objectivity in a domain D is equivalent to D s not being response-dependent. Response-dependence roughly amounts to the following:

RESPONSE-DEPENDENCE The concepts used to pick out D are such that their correct application to an object x is determined by how x affects us. (*cf.* Rosen (1994, 298))

A paradigm case is color-concepts: plausibly the concept of redness is a concept that is correctly applied only to objects that cause red-sensations in us (given that we are in normal environments, with normal lighting, have normal perceptual systems, etc.). This is all very rough, but it gives enough to illustrate one application of Rosen’s master argument.

Another response-dependent concept is, for Rosen, *constitutionality*. Roughly, laws are constitutional just in case the Supreme Court is disposed to judge them so; constitutionality is determined in part by how a law “affects” the Supreme Court. But this property of constitutionality, Rosen argues, does not show the notion to fail to be objective in the slightest sense. RESPONSE-DEPENDENCE fails as a characterization of non-objectivity since constitutionality satisfies it, yet is not thereby non-objective. Here is his argument:

So far we have been given no reason to think that the facts about what a certain group of people would think after a certain sort of investigation are anything but robustly objective. The facts about how the court would rule are facts of modal sociology ...but on the face of it they possess the same status as the facts about what any other collection of animals would do if prompted with certain stimuli, or set [*sic*] a certain problem. The facts about what the court would do with a given case ...are thus, for all we’ve said, features of the objective world. And if the facts [about constitutionality] just are these very facts, then [we have] no special grounds for thinking of them as less than entirely real. (Rosen 1994, 300)

Here Rosen is pointing out that the fact that constitutionality is response-dependent is itself a perfectly objective fact by any measure. It is just as objective as the fact that a mouse would run away if it saw a cat, or the fact that an electron would repel another nearby electron; all of these are dispositional facts grounded in fully objective psychological facts and other environmental factors. That is the import of the claim that “the facts about what the court would do with a given case ...are ...feature of the objective world”. This claim can be generalized: not only is the fact that laws affect the Court in certain ways an entirely objective fact,

an analogous claim is moreover true for *any* fact p and *any* response-dependent property. That is (where ‘O’ is the operator ‘it is objective that’ and ‘RD p ’ means ‘ p ’ is response-dependent):

O(RD p)

Of course is is objectivity of one kind of fact—the response-dependence of p —but we might still wonder about the objectivity of another kind of fact—namely the objectivity of p itself. Hence in order to conclude from this argument that p itself must be objective, that is

O p ,

Rosen needs the following premise:

O(RD p) \rightarrow O p

Rosen seems to accept the need for this principle in his argument and supports it with the following remark: “*Intuitively*, if the facts in the contested class can simply be read off in a mechanical way from the facts in an uncontroversially objective class, then there can be no grounds for denying the same status to facts in the contested area”.¹²

This sets up a master argument against the possibility of any non-objective facts. For given a broadly naturalistic framework, analogous premises will hold for *any* proposed characterization of objectivity Φ . Rosen can then substitute the proposed characterization in the following argument to show that any candidate for non-objectivity p is in fact fully objective:

The Master Argument

O(Φp)

O(Φp) \rightarrow O p

Therefore, O p

Let us call the second premise in The Master Argument the *collapse premise*:

COLLAPSE O(Φp) \rightarrow O p

¹²Rosen (1994, 301). Dummett appears to deny this:

The realist and the anti-realist may agree that it is an objective matter whether, in the case of any given statement of the class, the criteria we use for judging such a statement to be true are satisfied: the difference between them lies in the fact that, for the anti-realist, the truth of the statement can only consist in the satisfaction of these criteria, whereas, for the realist, the statement can be true even though we have no means of recognising it as true. (Dummett 1978, 147)

We can now ask whether COLLAPSE is indeed a structural feature of objectivity.

There is reason to doubt that it is, so long as we restrict our attention to epistemic conceptions of objectivity. Take BELIEF-KNOWLEDGE CONNECTION as a toy theory of non-objectivity. This means that if, necessarily, any claim p in a domain D is such that competent subjects who believe p thereby know p , then D is not an objective domain. If the collapse principle shows that objectivity is trivial, then for any putatively non-objective domain D^* , it must follow from the fact that it is objective that D^* satisfies the proposed criterion for non-objectivity that D^* is, in fact, fully objective.

Given this background, let us make some simplifying assumptions: Suppose first that we are in a world where every agent is logically and probabilistically competent. And moreover, let's assume that everyone in this world believes all of the propositions in a proposed non-objective domain D^* , and that they are not in an environment that is sufficiently friendly to allow for knowledge of any other proposition. Given that belief by a competent agent is sufficient for knowledge, as *Belief-Knowledge Connection* says, then on these simplifying assumptions, the claim that Ψ is non-objective implies that it is known. And, since the agents in question are, *ex hypothesi*, not in an environment that allows for knowledge of other objective domains, that something is known implies that it is objective. This gives us the following assumptions:

$$\neg O\Psi \rightarrow K\Psi$$

$$K\Psi \rightarrow \neg O\Psi$$

These imply:

$$\neg O\Psi \leftrightarrow K\Psi$$

Since on this proposal the claim that a domain is non-objective is just the claim that it is known, COLLAPSE is then (given these assumptions) just the claim that is obtained by replacing the schematic variable ' Φ ' with ' Kp '. (Recall that COLLAPSE is supposed to hold for any characterization of non-objectivity.) If p is a candidate non-objective claim, then it is on these assumptions known:

$$OKp \rightarrow Op$$

But since our simplifying assumptions make instances of ' $\neg O\Psi \leftrightarrow K\Psi$ ' holds for any claim, the occurrences of ' O ' in the collapse principle can be replaced with ' $\neg K$ '. Rosen's collapse principle is equivalent to the following:

$$\neg KKp \rightarrow \neg Kp$$

This however is just the KK principle (typically written as the equivalent ' $Kp \rightarrow \neg KKp$ ') which has compelling counterexamples.

The counterexamples to the KK principle do not disappear once we make the assumptions necessary to make the BELIEF-KNOWLEDGE CONNECTION characterization of objectivity entail $O\Psi \leftrightarrow \neg K\Psi$. The counterexamples to this principle in Williamson (2000, Ch. 4) have the following general structure: agents might come to believe p in a way that has whatever features necessary for knowing p , but not be in a position to come to believe in the same way that their belief in p has the knowledge-making features. In some cases, it may even be highly unlikely on one's evidence that one has formed a belief in p in a way that makes the belief constitute knowledge.¹³ Subjects who are logically and probabilistically competent and who are only in friendly epistemic circumstances with respect to p are not shielded from these kinds of situations. Mere formal coherence does not guarantee that, whenever one has formed a belief in a way that makes for knowledge, one also knows that one has in fact formed a belief in this way. And the friendliness of the circumstances for forming beliefs about p may not extend to beliefs about whether one knows p : one might not, after all, know that the circumstances are so inhospitable.

The point of the present section is that the COLLAPSE principle is not a structural feature of objectivity, as it is demonstrably false given some epistemic characterizations of the notion. Nothing in this point relies on the thought that BELIEF-KNOWLEDGE CONNECTION is the best epistemic characterization of the notion. But it is worth noting in closing that the COLLAPSE principle will likely be false for the same reasons given a range of more sophisticated epistemic characterizations of objectivity. First many instances of BKCS, will still permit the same kind of counterexamples: any conditions we substitute for the schematic variable C will be such that one can satisfy them and yet not know one satisfies them. And second even if we move away from the epistemic category of knowledge altogether COLLAPSE will still be highly questionable. For it is in general perfectly possible to be in a privileged epistemic position with respect to p while not being in a privileged epistemic position with respect to the fact that one is in a privileged epistemic position with respect to p . The fact that COLLAPSE ignores these distinctions, it should be rejected as a structural feature of the epistemic conception of objectivity.

¹³One example here is Williamson's "clock case" (Williamson 2011): you are staring at an unmarked clock, which as a matter of fact is such that its minute hand is pointing at 25 minutes past the hour. Your perceptual abilities are not good enough to know this, however; at best you can, for any point on the clock n , know that the minute hand is pointing to $n \pm 2$ minutes. Hence in your actual situation the strongest proposition you can know that it is exactly between 23 and 27 minutes past the hour. But you cannot know that you know this; in fact it is highly unlikely on your evidence that you do: since for all you know, the minute hand could be pointing towards 23, 24, 26, or 27 minutes past, and if it is, then the strongest proposition you can know is not that it is between 23 and 27 minutes past the hour.

4 Constructivism

Attempts to deny objectivity to a domain have some philosophical appeal. If successful, they provide one with an easy account of how we can come to know about the domain. Non-objective domains by definition remove the possibility of belief (under the right conditions) without knowledge, then knowledge about the domain requires only belief in the right conditions. Views which purport to deny objectivity to a domain—such as our Constructivist and Subjectivist examples from §1—might thereby have a crucial advantage over other ethical or meta-ethical views, if they can provide an uncontroversial and satisfactory account of moral knowledge while other views cannot.

This advantage has in fact been claimed for multiple versions of the Constructivist view. Below I will examine two instances of this claim before moving to some tentative generalizations about it. The framework for investigation is as follows: while it is common to assume that Constructivism is a meta-ethical view according to which ethics is not objective,¹⁴ there is (per our 1 discussion) no guarantee that this usage picks out a very natural, theoretically interesting kind. On the other hand, many instances of BKCS capture theoretically interesting notions of objectivity. Moreover, *if* Constructivism entailed these instances, its proponents will clearly have been correct in claiming that it secures epistemological advantages over its objective rivals. But it is not guaranteed that Constructivism will entail any such instances; merely claiming that one's view has epistemic benefits does not by itself secure these benefits. Hence ethics many, in the end, fail to be a non-objective view of ethics in any philosophically interesting sense.

I will proceed by first outlining the plausible instances of BKCS that some versions of Constructivism might plausibly be thought to entail. Then I argue that these entailments do not in fact hold. The immediate upshot of this is that some Constructivist views in the literature do not entail non-objectivity for ethics when we restrict ourselves to some philosophically interesting epistemic senses for the term. But some more general lessons may be in the offing, which I sketch. First, I outline an argument that the remaining problems in moral epistemology are, for the Constructivist, exactly the same as those faced by her realist counterparts. I conclude with a second observation, which is that the fact that the few versions of Constructivism I consider here fail to entail that ethics is non-objective may not be accidental. Instead, it may be a general feature of epistemic conceptions of objectivity that, given some plausible epistemological assumptions, objectivity is trivial after all.

¹⁴Recall that we are interested here in a distinctive epistemic sense of 'objective'. There may be other senses of the term on which Constructivism does (aim to) vindicate objectivity; see O'Neill (2003) for discussion.

4.1 Idealized Constructivism

As a first case study, we can focus on one particular development of the Constructivist view as given by Sharon Street (2008). Street formulates the view roughly as follows (I have made minor changes to keep the ideology in line with that of the rest of this paper):

CONSTRUCTIVISM The fact that A ought to ϕ is constituted by the fact that the judgment that A ought to ϕ withstands scrutiny from the standpoint of A 's other judgments about what A ought to do. (Street 2008, 223)

In other words, what makes it true that A ought to ϕ is the fact that, were A to reflect on her judgment that she ought to ϕ , she would not retract it in light of her other normative commitments—the judgment would *survive scrutiny*. I will call this *Idealized Constructivism*, since this view grounds facts about what one ought to do in facts about what a certain idealized agent, who has at least considered all of her relevant normative judgments, would judge that she ought to do.

Street claims that an Idealized Constructivist view has advantages in the epistemic domain since it does not imply that it is mysterious or coincidental that our ethical beliefs should constitute epistemic successes. This appears to be the crucial component of Street's "Darwinian Dilemma" against realism. She says on behalf of the Constructivist view:

The breaking of our bones *is* bad, in other words, and we're well aware of this. But the explanation [of why we are aware of this] is not that it is true independently of our attitudes that the breaking of our bones is bad and we were selected to be able to notice this; the explanation is rather that we were selected to take the breaking of our bones to be bad, and this evaluative judgement withstands scrutiny from the standpoint of our other evaluative judgements (to speak, for example, in the voice of the constructivist antirealist).¹⁵

Street in this passage identifies the source of a realist's difficulties in moral epistemology: picturesquely, there are supposed to be some "independent" moral truths, and then somehow, according to the realist, we come to know (become "aware") of these independent facts. Let's grant that the realist has, as Street says, an explanatory task that is difficult to meet here. Her Idealized Constructivist is supposed to avoid a difficulty of the same kind: since judgments under scrutiny *constitute* facts about what one ought to do, there are supposedly no epistemological mysteries surrounding the view.

Of course Street is here offering just a sketch of a point favorable to the Constructivist and not a full-fledged argument. But on the face of it these claims together do not remove all epistemological puzzles for the Idealized Constructivist. What Street has given us is an account of what makes beliefs about matters

¹⁵Street (2006, 154), see also Street (2008, §7).

normative *true*. And, on this account, the counterfactual facts about those very beliefs that make them true. But this is an account of what makes ethical beliefs true, and not how we can *know* them. True belief and knowledge are distinct, and we can ask below whether the Idealized Constructivist view has an advantage over realist views with respect to explanation of the latter, or whether she has merely explained why some of our normative beliefs are true.

The question is connected to the objectivity of ethics on the Idealized Constructivist view. *If* it follows from the view that normative claims are known in appropriate conditions, then the view entails that ethics is not objective in the sense of one instance of BKCS. Of course the Idealized Constructivist will want to insist that not everyone actually knows normative claims: some may fail to have knowledge in virtue of not reflecting on the relevant issues at all, and others may form false beliefs on the basis of faulty deductive or probabilistic reasoning. And the Idealized Constructivist would not want to claim that everyone has true beliefs about what their beliefs after scrutiny would be like. (While it is presumably easy to know what one would judge under scrutiny, it is not impossible for someone to be misled into wrongly thinking that they would judge under scrutiny that pain is good. Evolution does not shield us from gross mistakes of this kind, and the Idealized Constructivist should not be saddled with the claim that even these people succeed in having normative knowledge.) Thus the instance of BKCS we are interested in will state that *if* a subject avoids these errors, and avoid other errors of a similar kind, then she will thereby know normative claims if she believes them. This is a substantive notion of non-objectivity, and we can ask whether the Idealized Constructivist view entails it.

In more detail the proposal is that Idealized Constructivism entails CONSTRUCTIVIST BELIEF-KNOWLEDGE CONNECTION, which would appear to be a respectable notion of non-objectivity:

CONSTRUCTIVIST BELIEF-KNOWLEDGE CONNECTION (CBKC): Necessarily, for any subject s , if s logically and probabilistically competent, and knows all of the psychological facts about what s would judge in counterfactual circumstances where she considers whether s ought to ϕ , then if s believes that s ought to ϕ , then s thereby knows that s ought to ϕ .

This is one way of making sense of the claim that Idealized Constructivism has an advantage over its realist rivals in moral epistemology. That is, if it CBKC true, then all that is required is for one to have knowledge is acquire psychological knowledge via introspection about which of one's evaluative judgments one would endorse after scrutiny, plus other minimal conditions of competence. Since many subjects do meet these requirements, many subjects have ethical knowledge. Let us encode this as the "Idealized Constructivist Argument for Easy Ethical Knowledge", or *Easy Knowledge Argument* for short:

Easy Knowledge Argument

CBKC is true

It is easy to satisfy the antecedent of CBKC

Therefore, it is easy to have normative knowledge

The Constructivist would appear to be warranted in assuming the second premise, and at any rate I will not contest it here. A Belief-knowledge connection thesis that is restricted to subjects who have basic knowledge of their own psychology and meet other minimal conditions on competence does not thereby make normative knowledge difficult—it if is true, then many actual people will have normative knowledge.

The remaining question for the Easy Knowledge Argument is whether CBKC follows from Idealized Constructivism. I will argue that it does not: even if we assume that Idealized Constructivism is true, I will claim, CBKC is false.

Suppose a is an agent who satisfies the antecedent of CBKC. She is logically and probabilistically competent and she knows all of the relevant psychological facts about herself. This last condition includes the claim that she that if she were to reflect on her evaluative commitments, she would judge that she ought to ϕ . That is:

$KJ(a \text{ ought to } \phi)$

(Here K and J are relativized to the agent a , and J is read as ' a would, if she were to scrutinize her evaluative commitments, judge that . . .'.) We can also assume, as Street claims, that it is an evolutionary advantage for people like a to believe that they ought to ϕ , and a in fact believes this:

$J(a \text{ ought to } \phi)$

Hence a satisfies all of the conditions of CBKC. It is then true if it follows from these assumptions that she knows she ought to ϕ , that is

$K(a \text{ ought to } \phi)$.

But this latter claim only follows given the additional premise that she knows that if a judges that she ought to ϕ , then she ought to ϕ . The following argument is valid, given a plausible closure principle according to which the entailments of known premises are also known. (We needn't worry about the subject failing to believe the relevant conclusion because we are assuming that the subject is deductively competent, and hence believes the entailment on the basis of the relevant deduction.) We can call this a *Closure argument*:

$KJ(a \text{ ought to } \phi)$

$K(J(a \text{ ought to } \phi) \rightarrow (a \text{ ought to } \phi))$

Therefore, $K(a \text{ ought to } \phi)$

Unfortunately the Idealized Constructivist view does not entail the second premise. The Idealized Constructivist view supports the conclusion that the conditional embedded in the premise is *true*: given that a subject would judge under scrutiny that she ought to ϕ , the view on offer just is that it is true that she ought to ϕ . But even if this conditional is true, that is different from the conditional in question being *known*. Once we replace the second premise with a claim the Idealized Constructivist is entitled to, we have an invalid argument:

$KJ(a \text{ ought to } \phi)$

$J(a \text{ ought to } \phi) \rightarrow (a \text{ ought to } \phi)$

Therefore, $K(a \text{ ought to } \phi)$

The simple reason why the second argument is invalid is that the agent in question might not have any beliefs at all concerning the second premise, and hence would not in any way have epistemic access to her obligations merely in virtue of the first premise being true: a belief that a ought to ϕ would, under these circumstances, amount to nothing more than a mere guess.

The reason why the Idealized Constructivist cannot help herself to the first argument, which we have granted is sound given a plausible closure principle, is as follows. Not only does the view not guarantee that the subject has beliefs concerning the conditional which, if absent, would preclude knowledge (since belief is necessary for knowledge). In addition the Idealized Constructivist view is consistent with subjects who have formed beliefs on the matter falling short of knowledge. Suppose a actually does believe the embedded conditional, and hence has a true belief. Still she might not know it. Even if she actually holds a true belief that $J(a \text{ ought to } \phi) \rightarrow (a \text{ ought to } \phi)$, she might have *disbelieved* the conditional in nearby worlds, even if she retains knowledge of the relevant psychological facts and remains logically competent in these worlds. Subjects could easily have believed that their normative judgments formed under scrutiny would still be false. That is, the following could easily have been true:

$B\neg(J(a \text{ ought to } \phi) \rightarrow (a \text{ ought to } \phi))$.

Given some very plausible constraints on knowledge, the nearby possibility of this kind of error is incompatible with knowledge in the actual world. (Indeed the need for this kind of luck seems to be at the heart of Street's criticism of the epistemology of realist views.) Hence there is a natural argument from an assumption of the Idealized Constructivist view plus the conditions specified in the antecedent of CBKC to the conclusion

$\neg K(J(a \text{ ought to } \phi) \rightarrow (a \text{ ought to } \phi))$.

The mere truth of the embedded conditional is not enough to extend *a*'s knowledge to the normative realm. CBKC is false given Idealized Constructivism.

This is a quite general point—even if there is a connection between easy-to-know counterfactual judgment and ethical truth, we might wrongly have believed that there was no such connection, and hence fail to know that the connection holds. After considering a second case study, I will say more about the significance of this argument in closing remarks in §5.

But before moving on it deserves emphasis here that the point is not to argue that entailments such as

$$J(a \text{ ought to } \phi) \rightarrow (a \text{ ought to } \phi),$$

unknowable even if true. Upholding a principle that yields such a claim would lead to an untenable skepticism over philosophical claims. Rather the point is that there is nothing special in the Idealized Constructivist view that implies that CBKC. It might still be *true* that one can know the relevant conditionals, and hence that claims about obligation are knowable. But this is not a consequence of the Idealized Constructivist's view of what constitutes moral facts alone. Hence one does not simply get to satisfy the antecedent of CBKC and *thereby* come to know normative facts just by believing them. Additional facts will need to be cited in explaining knowledge. Hence rejecting skepticism does not re-open the case in favor of CBKC. And moreover the additional explanatory machinery may be something the realist can help herself to as well—a possibility I explore more in §5.

4.2 *Relativism*

The previous subsection argued that Idealized Constructivism fails to entail that ethical knowledge is easy, and hence it does not entail that ethics fails to be objective in the epistemic sense of the term. This is not an isolated feature; other views in the same spirit have analogous consequences. Here I briefly consider a “Relativist” implementation of the Constructivist idea and show how similar claims about moral epistemology have been made on its behalf. These fail for the same structural reasons as before.

I will use the name *Relativism* for a version of Constructivism that makes use of a development in the philosophy of language of the the notion of assessor- or perspective-sensitive truth. The basic idea is that the same claim may be true relative to one context of assessment (or “perspective”), but false relative to another. Sentences involving the predicate ‘tasty’ are one candidate example of sentences that express assessor-sensitive claims: when it I say ‘rotting flesh is not tasty’, what I say is true as assessed from the standpoint of another person, while the same claim is false as assessed from the perspective of a vulture.¹⁶

¹⁶Cf. Cappelen & Hawthorne (2009).

As has been noted in the literature on relativism, this approach requires a notion of “relative truth”—truth from a perspective or context of assessment. In what follows I will assume that this is a theoretical notion in the toolkit of the relativist that is distinct from our ordinary notion of truth. The ordinary notion obeys disquotational principles: from an utterance of ‘S is true’ one can validly infer ‘S’. But with a relativized truth predicate—call it $true_R$ —a similar principle does not apply. For we can recognize that a sentence is true from some standpoint which is not ours and so accept that a claim about relative truth without accepting the claim itself. Someone who knew the dietary habits of vultures might accept ‘rotting flesh is tasty’ is $true_R$ (for the vulture) while rejecting ‘rotting flesh is tasty’.¹⁷

Relativism, then, is the view that normative claims about an individual are made true (in the ordinary sense) by what is $true_R$ from the standpoint of that individual. More precisely:

RELATIVISM The fact that A ought to ϕ is constituted by the fact that the claim that A ought to ϕ is $true_R$ from A ’s standpoint.

This is a position elaborated and defended in Schafer (2014). As with Street I will not focus on the details of the view but will instead just focus on the epistemological benefits that Schafer claims for Relativism. He first puts forward the principle below as a plausible claim about warranted judgment for the Relativist:

KNOWLEDGE-RELATIVE The judgment/assertion that P is warranted just in case the judge/asserter is a position to know that P is $true_R$ relative to his normative perspective. (Schafer 2014, 89)

And he goes on to say that this makes knowledge of ethical claims easy and unmysterious according to the Relativist view:

And it is not difficult to explain how we might generally be in a position to satisfy these norms. In this way, the perspectivalist can explain why normative judgments and assertions should be relatively easy to make in a warranted fashion without taking on the burdens involved in providing a non-skeptical realist epistemology. And so, if it becomes apparent that the project of providing such an epistemology is doomed, the perspectivalist is well placed to account for our practice of normative judgment and assertion by other means. (Schafer 2014, 90)

Schafer’s **KNOWLEDGE-RELATIVE** principle is stated in terms of warranted judgment instead of knowledge, and for good reason: while there is some sense in acknowledging that the judgment of someone with a different perspective is

¹⁷Schafer (2014, XX).

warranted, it is incoherent, if you disagree, to allow that it is knowledge. We might admit that the vulture is warranted in judging that rotting flesh is tasty, since it is true from his perspective. But we can't admit that the vulture *knows* that rotting flesh is tasty, as it would follow that rotting flesh *is* tasty. The relativist framework is supposed to allow us to acknowledge the truth of the vulture's judgment from his own perspective while retaining our own judgment at the same time that rotting flesh is not tasty, but granting knowledge to the vulture will not allow this.

Without trying to make sense of "warranted judgment" that plays the role needed by KNOWLEDGE-RELATIVE, we can focus on an entailment of the principle that ought to hold if it is to be a substantive and interesting thesis. In particular, the notion of "warrant" in the original KNOWLEDGE-RELATIVE should have a tight connection to claims about knowledge so long as the principle is assessed from one's own perspective. We can't uniformly grant that judgments that are both warranted for other speakers and are true from their standpoint are thereby known. However, once we restrict the judgments at issue in KNOWLEDGE-RELATIVE to judgments about what is true from the perspective of our own perspective, then given KNOWLEDGE-RELATIVE, these judgments will be both true *and* warranted. If the notion of warrant is an epistemologically interesting one, then it should follow that these judgments are also knowledge.

Thus any speaker should accept the following instance of BKCS, where c is their own standpoint:

RELATIVIST BELIEF-KNOWLEDGE CONNECTION (RBKC) Necessarily, for any subject s who shares standpoint c , if s knows that it is true_R from c that s ought to ϕ , then if s believes that s ought to ϕ , then s knows that s ought to ϕ .

But RBKC does not follow from Relativism, for reasons that are analogous to the reasons for which CBKC does not follow from Idealized Constructivism. Hence the Relativist has not show that KNOWLEDGE-RELATIVE is true on her view.

Let a be a subject who shares our evaluative perspective, c . In addition to being logically and probabilistically competent, we can suppose that a knows that a normative claim about herself is true from her standpoint. That is, it is true_R relative to c that a ought to ϕ . Since c is, *ex hypothesi*, both a 's standpoint and ours, it is not only true_R that a ought to ϕ , it is true simpliciter. If RBKC is true, it follows that if a believes this truth they also know it: that is, a knows that a ought to ϕ .

But there is a variation on the argument from the last subsection that this principle is false even given RELATIVISM. One way for the knowledge of truth_R to secure the conclusion would be via the following Closure argument, where TR_c abbreviates true_R from our shared standpoint c :

$$K(TR_c(a \text{ ought to } \phi))$$

$$K((TR_c(a \text{ ought to } \phi)) \rightarrow (a \text{ ought to } \phi))$$

Therefore, $K(a \text{ ought to } \phi)$

This argument is sound given a plausible closure principle, but its second premise is not one that follows from the Relativist's statement of her own view. While it may be easy to know what is true_R from one's own perspective, truth simpliciter is a distinct notion, and knowledge of truth_R does not ipso facto amount to knowledge of truth. And in fact there are plausible arguments from structural features of knowledge that the conditional is not in fact known on the Relativist view. It would seem that one could easily have believed that truth_R from one's perspective and truth simpliciter about obligation come apart; one, for instance, could easily have believed that one's evaluative perspective is mistaken (indeed many non-Relativists actually have such beliefs). In such worlds the following holds:

$$B \neg ((TR_c(a \text{ ought to } \phi) \rightarrow (a \text{ ought to } \phi))).$$

And if knowledge requires freedom from risk of such errors, then it follows that the second premise is false.

What the Relativist view does guarantee, in virtue of its claim that what is true_R relative to a constitutes facts about a 's obligations, is that the embedded conditional in the above argument is true:

$$K(TR_c(a \text{ ought to } \phi))$$
$$TR_c(a \text{ ought to } \phi) \rightarrow (a \text{ ought to } \phi)$$

Therefore, $K(a \text{ ought to } \phi)$

But this argument is invalid. The truth of the premises is quite compatible with one's failing to believe anything whatsoever concerning one's obligations, or being at risk of inferring the negation of the claim the conclusion says one knows.

As before, this is not to argue for skepticism about knowledge of obligation on the Relativist view. The only argument here is that the Relativist view (like the Idealized Constructivist) does not by itself entail that ethical knowledge is easy to come by; by satisfying the conditions in RBKC one does not *thereby* come to have ethical knowledge, and non-objectivity as captured by RBKC does not follow from RELATIVISM. Earlier I suggested that the additional epistemological machinery that might be employed to avoid skepticism for the Constructivist might be something the realist can employ as well as an aid in moral epistemology. I expand on the possibility of (and potential significance of) epistemological parity between these views in the concluding section.

5 Conclusion: triviality after all?

In closing I will draw out two tentative conclusions from the preceding section. First, the failure of these views to entail that ethical knowledge is easy plausibly

shows that these views face the *same* epistemological questions as the realist views they are meant to replace. And second, the foregoing together with the methodological guidelines laid down in §1 raise the prospect of a pessimistic conclusion concerning the central question of this paper, namely the conclusion that objectivity is in fact trivial.

5.1 *Structural parity*

We concluded §4 with the thought that, even if Relativism and Constructivism are true, the relevant instances of BKCS are not. But we were careful to dissociate this claim with the skeptical claim that moral knowledge is impossible given constructivism: instead the limited conclusion here is just that the Constructivist's view by itself does not imply that normative knowledge is easy to come by. Of course this leaves open the possibility that the proponents of such a view can appeal to some epistemological principles to explain why ethical knowledge is possible. However the threat looms that the additional principles will also be helpful to the realist, and thereby deny the Constructivist her claimed advantage over matters epistemic.

In more detail the situation is as follows. Recall that the Constructivist has a view on which (i) it is easy for agents to know what they would judge under scrutiny, and (ii) what these agents know constitutes the facts about what they ought to do. But these claims do not entail the second premise in the following Closure argument:

$KJ(a \text{ ought to } \phi)$

$K(J(a \text{ ought to } \phi) \rightarrow (a \text{ ought to } \phi))$

Therefore, $K(a \text{ ought to } \phi)$

Instead, they entail the conditional

$J(a \text{ ought to } \phi) \rightarrow (a \text{ ought to } \phi)$

which, when substituted for the second premise, makes the argument invalid.

Given this set-up, there is a natural way of seeing a realist as saddled with an analogous problem. The realist will have a view on which a certain natural property—say, the maximization of happiness, or *MH* for short, constitutes obligation. It will be easy to know, given a general non-skeptical outlook, which actions have this property. That is:

$K(MH\phi)$.

Of course the target conclusion for a realist view will be that the agent in question knows her obligations, that is:

$K(a \text{ ought to } \phi)$.

This puts us in a familiar situation: on the one hand, there is a Closure argument which would secure the needed conclusion as follows:

$K(MH\phi)$

$K(MH\phi \rightarrow (a \text{ ought to } \phi)),$

Therefore, $K(a \text{ ought to } \phi)$

However familiar concerns will arise for the claim that the second premise is true. The realist is instead is only entitled to the truth of the embedded conditional, but replacing the second premise with this conditional renders the Closure argument invalid:

$MH\phi \rightarrow (a \text{ ought to } \phi).$

The Constructivist's view of the metaphysics of normative facts doesn't by itself explain how normative knowledge is easy; at best her view gives an especially clean story about how one can come to know the facts that *constitute* normative facts. There may well be additional principles which explain how this knowledge can be transformed into normative knowledge. Potentially these principles would work by explaining how a "connecting conditional"—a conditional connecting non-normative claims to normative claims—is not only true but also known. The prospect then remains that the realist's connecting conditionals will also be knowable for exactly the same reasons. If so, the Constructivist's epistemological advantages disappear.

5.2 *An argument for triviality*

This paper opened by asking the question, 'is objectivity trivial?' and rejected one argument for this conclusion in §3. The potential for parity between Constructivism and realism suggests a general argument bearing on this question.

Recall that, if objectivity in the epistemic sense is trivial, then it follows from minimal assumptions that, for any domain D whatsoever, the facts in D are objective. We have been focusing mainly on the domain of normative facts, but if Constructivism and other anti-realist views of the normative fail to satisfy a viable form of non-objectivity, then a pessimistic conclusion concerning the possibility of non-objectivity in any domain looms. We argued in §4 that some especially natural instances of BKCS fail to follow from various versions of Constructivism. The question is, are there any *other* viable notions of non-objectivity which take the form of an instance of BKCS and which *do* follow from the Constructivist view?

One suggestion that is to build conditions that directly entail knowledge of ethical facts into the conditions in an instance of BKCS. For instance the Idealized Constructivist could accept an instance that applies only to agents who both (i) know what they would judge under scrutiny, and (ii) know that, if they were to judge that they ought to ϕ under scrutiny, then it is true that they ought to ϕ . In

other words, we build into the conditions for application of the instance that the relevant agents know that Idealized Constructivism is true. This is the following instance:

STRENGTHENED CONSTRUCTIVIST BELIEF-KNOWLEDGE CONNECTION (SCBKS) Necessarily, for any subject s , if s is logically and probabilistically competent, and s knows all of the psychological facts about what s would judge in counterfactual circumstances where she considers whether s ought to ϕ , and knows that if she knows that s would judge that s ought to ϕ under scrutiny, then s ought to ϕ , then if s believes that s ought to ϕ , then s thereby knows that s ought to ϕ .

It should be fairly uncontroversial that Idealized Constructivism entails SCBKS. But this does little to help with formulation of a substantive notion of objectivity. Analogous Belief-Knowledge connections will hold for *any* view of any subject-matter whatsoever. For instance, the standard realist view one which obligation is constituted by happiness-maximization would entail an analogous instance of BKCS according to which any competent agent who knows that ϕ -ing maximizes happiness, and knows that if ϕ -ing maximizes happiness, then one is obligated to ϕ , is such that, if she believes that she ought to ϕ , then she thereby knows that she ought to ϕ .

Thus the most straightforward way to obtain instances of BKCS that do follow from prima facie non-objective views turn out not to capture a legitimate notion of non-objectivity. This is not of course a decisive argument that objectivity must be trivial; we could only arrive at this conclusion by considering all candidate notions of non-objectivity and showing of each that it is not entailed by the relevant views. But here one of the methodological guidelines laid down at the outset is relevant: genuine candidate notions of objectivity should be natural theoretical kinds, and not highly gerrymandered. Thus the search for an instance of BKCS that does in fact follow from Constructivism (unlike CBKC) but for which no analogue is available to the realist (unlike SCBKC) will be highly constrained. The COLLAPSE principle failed to yield a sound argument for the triviality of objectivity, the existence of a substantive epistemic conception of objectivity may be a chimera after all.

References

- Cappelen, H. & Hawthorne, J. (2009), *Relativism and Monadic Truth*, Oxford University Press.
- Dummett, M. (1963), Realism, in 'Truth and Other Enigmas', Duckworth.
- Dummett, M. (1978), *Truth and Other Enigmas*, Harvard University Press.
- Dworkin, R. (1996), 'Objectivity and Truth: You'd Better Believe It', *Philosophy and Public Affairs* 25(2), 87–139.

- Harman, G. (1975), 'Moral Relativism Defended', *The Philosophical Review* **84**, 3–22.
- Hawthorne, J. & Srinivasan, A. (2013), Disagreement without transparency: Some bleak thoughts, in D. Christensen & J. Lackey, eds, 'The Epistemology of Disagreement', Oxford University Press.
- Kölbel, M. (2002), *Truth without Objectivity*, Routledge.
- Kölbel, M. (2004), 'Faultless Disagreement', *Proceedings of the Aristotelian Society* **104**(1), 53–73.
- MacFarlane, J. (2007), 'Relativism and Disagreement', *Philosophical Studies* **132**, 17–31.
- Nagel, T. (1989), *The View from Nowhere*, Oxford University Press.
- O'Neill, O. (2003), Constructivism in Rawls and Kant, in S. R. Freeman, ed., 'The Cambridge Companion to Rawls', Cambridge University Press.
- Putnam, H. (1987), *The Many Faces of Realism*, Open Court.
- Railton, P. (1986), 'Moral Realism', *The Philosophical Review* **95**(2), 163–207.
- Rosen, G. (1994), Objectivity and Modern Idealism: What is the Question?, in M. Michael & J. O'Leary-Hawthorne, eds, 'Philosophy in Mind: The Place of Philosophy in the Study of Mind', Kluwer Academic, pp. 277–319.
- Sainsbury, R. M. (1996), 'Crispin Wright *Truth and Objectivity*', *Philosophy and Phenomenological Research* **56**(4), 899–904.
- Schafer, K. (2011), 'Faultless Disagreement and Aesthetic Realism', *Philosophy and Phenomenological Research* **82**(2), 265–286.
- Schafer, K. (2014), 'Constructivism and Three Forms of Perspective-Dependence in Metaethics', *Philosophy and Phenomenological Research* **89**(1), 68–101.
- Stevenson, C. (1937), 'The Emotive Meaning of Ethical Terms', *Mind* **46**, 14–31.
- Street, S. (2006), 'A Darwinian Dilemma for Realist Theories of Value', *Philosophical Studies* **127**(1), 109–166.
- Street, S. (2008), Constructivism about Reasons, in R. Shafer-Landau, ed., 'Oxford Studies in Metaethics, vol. 3'.
- Williamson, T. (2000), *Knowledge and Its Limits*, Oxford University Press.
- Williamson, T. (2011), Improbable Knowing, in T. Dougherty, ed., 'Evidentialism and Its Discontents', Oxford University Press.
- Wright, C. (1992), *Truth and Objectivity*, Harvard University Press.
- Wright, C. (2001), 'On Being in a Quandary', *Mind* **110**, 45–99.