

Companions in Guilt Arguments and Moore's Paradox¹

Michael Campbell

Abstract: In a series of articles Christopher Cowie has provided what he calls a 'Master Argument' against the Companions in Guilt (CG) defence of moral objectivity. In what follows I defend the CG strategy against Cowie. I show, firstly, that epistemic judgements are relevantly similar to moral judgements, and secondly, that it is not possible coherently to deny the existence of irreducible and categorically normative epistemic reasons. My argument for the second of these claims exploits an analogy between the thesis that epistemic norms are non-categorical and G.E. Moore's paradox concerning first personal belief ascriptions. I argue that the absurdity of the assertion "I have evidence that *p* but no reason to believe it" shows that the norms of belief are categorical. I then consider the counter-argument that this categoricity is a 'conceptual' rather than an 'objective' requirement. By drawing on the work of Hilary Putnam and Charles Travis, I show that this counter-argument is unsuccessful. Putnam is one of the original proponents of the Companions in Guilt strategy. Thus, by supporting the CG argument through appeal to other Putnamian theses, I show that its insights can only fully be appreciated in the context of broader metaphysical and semantic lessons.

Keywords: Companions in Guilt, epistemic reasons, error theory, meta-ethics, moral objectivity, normativity.

1. Introduction

The Companions in Guilt (CG) strategy seeks to rehabilitate some supposedly problematic class of judgements by showing how its claims are relevantly similar to another, less contentious area of discourse (Lillehammer 2007). For an area of discourse to be problematic it must simultaneously lay claim to a standing and yet fall short of it. Let the area of discourse in question be morality and the standing in question be objectivity. Moral realists aim to show that – appearances to the contrary – moral judgements can be objectively true. Moral irrealists, by contrast, deny this in one of two ways. On one flank are expressivists, who re-evaluate the commitments of moral discourse so that moral judgements do not in fact aim to reflect some fact about stance-independent reality. On this view, since moral judgements have as their function

¹ This publication was supported within the project of Operational Programme Research, Development and Education (OP VVV/OP RDE), "Centre for Ethics as Study in Human Value," registration No. CZ.02.1.01/0.0/0.0/15_003/0000425, co-financed by the European Regional Development Fund and the state budget of the Czech Republic.

the expression of a subject's mental state, they involve no commitment to objective truth and thus embody no error.² On the other flank are error theorists, who take the objective purport of moral claims at face value, but insist that the world simply cannot be as it would have to be for a moral judgement to be true. On this view, moral claims are necessarily uniformly false.

The plausibility of each of these views turns crucially on one's understanding of both the relevant subject matter (the nature of 'morality') and the standard (what it is for something to be 'objective'). For those with a very permissive conception of objectivity, morality could only be debarred from meeting its demands if it were shown to be highly peculiar. On the other hand, for those with a very demanding conception of objectivity, morality will likely fail to meet its standards, irrespective of how it is understood.³ The irrealist task thus gets easier as our conception of objectivity becomes more demanding, and *vice versa*.

A stance must be taken on these issues at the outset, even if it is only provisional. This is problematic, however, as it is very hard to define our terms in ways that do not load the dice in favour of one side or the other. Later on I will argue that the CG strategy works by correcting some persistent distortions in our conception of objectivity. For this reason we cannot begin with a definition of the concept which is satisfactory to all parties. As a generally acceptable gloss on the concept we must confine ourselves to the following: for a claim to be objective, it must disclose some aspect of the world as it is in itself, rather than merely as it appears to the subject. Let's take morality to be a form of thought, the truth of the claims of which depends on the existence of categorically binding and action guiding reasons.⁴ Then, for morality to count as objective moral judgements must be capable of expressing a true thought concerning some feature of the world while at the same time providing suitably placed agents with categorical reasons for action. Of course, what it means to 'express a true thought' can be given a more or less demanding interpretation.⁵ For dialectical purposes we may begin by ceding to the irrealist the right to interpret this condition however they want, on the proviso that they recognize another area of discourse which includes judgements that are objective in this sense. This proviso leaves the space that the CG strategy exploits. By showing the companion

² On this understanding of expressivism see (Schroeder 2008).

³ This explains why error theorists, with their highly demanding conception of what it takes for a judgement to be objectively true, generally spend little effort characterizing the nature of morality. Given their assumptions, any discourse which entails the existence of categorical action-guiding reasons will *ipso facto* be debarred from objectivity.

⁴ This is not of course intended to be complete as a definition, but it is sufficient to motivate the irrealist position. See Cuneo (2007) and Joyce (2001) for related discussions.

⁵ See Blackburn (1993, ch.6) and Dreier (2004) for discussion of some of the issues here.

area of discourse to be relevantly similar to morality, we can thereby rehabilitate the objectivity of morality.⁶

2. Cowie's 'Master Argument'

The CG argument is a method for defending morality against construal in irrealist terms. It works by showing that the rejection of moral objectivity would incur an unacceptable theoretical cost; that the same considerations which purport to justify such a rejection would entail the rejection of objectivity in some other, apparently unproblematic, context. From here on I will follow standard procedure and take the other context in question to be the class of epistemic judgements. The CG theorist then reasons as follows: "It is implausible to hold either that epistemic reasons are merely subjective, or that judgements making normative epistemic claims necessarily fall short of objective truth. Therefore, since we tolerate the existence of objective epistemic reasons, we ought to tolerate the existence of objective moral reasons also."⁷

In a series of articles (2014, 2016) Christopher Cowie has argued that the CG defence cannot work.⁸ Cowie paints the following picture of the landscape. The error-theoretic argument runs as follows:

Conceptual premise: Moral judgements presuppose for their truth the existence of categorical reasons for action, where for a reason to be categorical is for it to obtain independently of contingently held desires or mere convention.

Metaphysical premise: Reasons exist only where there is some desire or social convention to explain their existence. (2016, 116)

Conclusion: Therefore, moral judgements cannot be true.

Against this, the CG theorist provides the following counter-argument:

Parity Premise: The metaphysical premise of the moral error-theoretic argument entails that there are no categorical epistemic reasons for belief.

Epistemic Existence Premise: There are some categorically normative epistemic reasons for belief. (2016, 116)

Conclusion: Therefore, the metaphysical premise of the error-theoretic argument is false.⁹

⁶ On this understanding, it is not sufficient for a judgement to count as objective that it be such that it would be converged upon by ideal deliberators. Objectivity is a matter of a judgement's revealing some aspect of the world to us. The notion of a judgement's being world disclosing should be understood as neutral between those like Cuneo (2001) who take moral realism to require that moral properties be mind-independent and those like McDowell (1985) for whom moral judgements can be objective even if moral properties are mind-dependent.

⁷ The term 'companion in guilt' was coined (I believe) by Putnam in his (1982). The most influential recent statement of the approach is Cuneo (2007).

⁸ In my exposition of Cowie I will confine myself largely to the refined version of the argument presented in his later article (2016).

Cowie thinks that the underlying logic of this defensive maneuver is fundamentally flawed. He thus believes not only that its extant instances are unpersuasive, but that the approach as a whole simply cannot yield a convincing defence of moral objectivity. His justification of this striking view is what he dubs his 'Master Argument' against the CG strategy. Its centerpiece is the claim that the parity and existence premises stand in tension with one another, such that any position that established the truth of the parity premise would permit rejection of the existence premise, and *vice versa*. Cowie aims to demonstrate this by showing that a CG argument involving epistemic reasons succumbs to this dilemma. He then takes it that this result will generalise to any other potential companion in guilt. His Master Argument runs as follows. (The example is my own.)

Suppose that walking through the forest Abe discovers a distinctively shaped paw print and judges on the basis of this that there is a tiger nearby. The paw print is then (being taken as) a piece of evidence for the proposition that there is a tiger nearby. Let an 'evidential support relation' (ESR) be a probability raising relation that holds between some given piece of evidence *e* and a given proposition, *p*.¹⁰ (2016, 117) How does the fact that *e* makes *p* more likely relate to Abe's reasons to believe that *p*? Cowie constructs the following dilemma:

Either categorically normative epistemic reasons for belief *are* evidential support relations or they are not. If they are, then the parity premise is false. If they are not, then the epistemic existence premise isn't established. (2016, 127)

Suppose that categorical normative epistemic reasons for belief are identical with evidential support relations. Why does this render the parity premise false? Because, Cowie holds, this identity enables a reduction of epistemic reasons to 'non-normative (i.e. 'descriptive') facts or properties' (9), when no such reduction is available in the moral case. Supposing this to be feasible, then the considerations that support the error theory's metaphysical premise do not generalise to the epistemic context. And so:

epistemic reasons may be categorical without being metaphysically and epistemologically problematic. This would warrant rejection of the parity premise. (2016, 123)

This yields a mixed view, a combination of epistemic realism and moral irrealism. On this view, we are to admit the existence of objective and categorical reasons, but to deny that any moral reasons are to be found amongst them.

On the other hand, suppose that categorical normative epistemic reasons for belief are not (identical with) evidential support relations. In this case, the

⁹ His exposition here follows that in Cuneo (2007).

¹⁰ For ease of exposition I have substituted 'proposition' (represented by the dummy variable *p*) for Cowie's 'hypothesis' (and his dummy variable *h*).

epistemic existence premise is undermined, according to Cowie, because the error theorist can simply bite the bullet and deny the existence of categorically normative reasons for belief. This means granting that probability raising relations hold between states of affairs, but denying that such facts in themselves entail any conclusions about a subject's reasons for belief. Abe has a reason to form beliefs on the basis of evidence only given her possession of an independent and extra-rational desire to attach credence to beliefs in proportion to their likelihood. This yields a denial of categorical reasons across the board; epistemic and moral irrationalism.

This uniformly sceptical position denies the epistemic existence premise of the CG argument. This premise is a conjunction of an existence claim and a conceptual claim. The former is the claim that there are (or must be) some epistemic reasons. The latter is that these reasons are (or must be) categorical. There are therefore two strategies open to the epistemic irrealist. The first of these is to deny that there are any epistemic reasons. This is *epistemic nihilism*. The second is to accept that there are epistemic reasons, but to maintain that they bind on us only hypothetically (in Kant's sense of that term).¹¹ This is *epistemic subjectivism*. For the sake of argument I will follow Cowie's lead and assume that epistemic nihilism is not a going concern.¹² Epistemic irrationalism is then equivalent to epistemic subjectivism.

This leaves us, then with two positions which can be adopted against the CG theorist. On the one hand, there is epistemic realism buttressed by the provision of a non-sceptical reduction of epistemic reasons to non-normative facts. On the other hand, there is epistemic irrationalism, in the form of a construal of epistemic reasons as binding only given a subject's interests. Cowie claims that the considerations which support the existence premise with respect to epistemic reasons undermine the parity premise, and *vice versa*. Since the only considerations that support epistemic realism ruin the putative analogy between the epistemic and the moral, the CG defence cannot succeed.

Cowie's Master Argument nicely brings out how daunting the CG theorist's task is. In order for a CG maneuver to be effective one must show both that judgements in moral and epistemic contexts are relatively similar, and that scepticism regarding the latter is not a going concern. Keeping this many balls in the air at once requires significant dexterity.

In order to undertake such a defence, I will argue as follows. Firstly, if the error theorist admits the existence of evidential support relations, she is thereby committed to the existence of categorical epistemic reasons.¹³ Thus, the

¹¹ This mirrors Cuneo's distinction between epistemic irrealists who deny the content and authority platitudes of epistemic norms, respectively (2007, 36-9).

¹² But see Streumer (2013).

¹³ For a related approach see (Evers 2015). Evers notes the possibility of giving a 'non-hypothetical reading' of epistemic conditionals, but does not advance any arguments in favour of such an account. What follows can be understood as an attempt to fill this lacuna.

epistemic existence premise is true. Secondly, like moral reasons, epistemic reasons cannot be reduced any collection of non-normative facts or properties. Thus, the parity premise is true. So, the CG strategy succeeds after all. I will take these points in turn, but with particular focus on establishing the epistemic existence premise. My defence here takes the form of a negative argument: I show that the epistemic existence premise cannot coherently be denied, because we can neither deny that there are epistemic reasons, nor understand such reasons as only hypothetically binding. I then show that the apparent plausibility of epistemic irrationalism depends on a mistaken thesis about the relation between mind and world, which (following Putnam) I dub *metaphysical realism*. I leave discussion of the parity premise to §8. There, I argue that both moral and epistemic reasons may (in some sense) be 'grounded in' non-normative facts, but we have no reason to think in either case that it is possible to give a non-sceptical reduction of reasons to their non-normative bases.

My strategy here closely follows Cuneo's, whose pioneering work a decade ago first popularised the CG strategy (Cuneo 2007). However, there are two important differences between my approach and his. Firstly, I emphasise more strongly the role of the Moorean paradox in showing the untenability of epistemic irrationalism. Secondly, I tie the defence of moral objectivity to broader metaphysical issues concerning the nature of objectivity. These two points are connected. Cuneo is limited in the use that he can make of the Moorean paradox precisely because he is committed to metaphysical realism. Witness his definition of the moral realist view that he defends:

It is in virtue of a fact's sharing the same 'informational content' as a given proposition that that fact guarantees that the proposition in question is true. Or to put the point in an admittedly more stilted manner: According to the alethic realist's position, for any true proposition that *p*, what the proposition that *p* is a proposition *that* is the same as what the fact *that p* is a fact *that*. It is this phenomenon of proposition and fact sharing the same content that ensures that there is the right sort of satisfaction relation between a particular proposition and fact. (Cuneo 2007, 28)

Although Cuneo is agnostic as to whether the relation between true propositions and facts should be understood in terms of a relation of correspondence, he is clear that propositions are made true solely by virtue of their standing in a particular relation (viz., 'sharing the same content') with the world. This is metaphysical realism in essence. As a result of his commitment to this view he is forced to see the Moorean paradox as arising from the commitments of the speech-act of assertion (Cuneo 2007, 43). Since there is no contradiction in A's believing that *p* when *p* is false, if there is a contradiction in A's asserting the conjunction of these facts then that must be a product of the rules of assertion.

This view prompts the objection that, by arguing from the commitments of our speech acts to the existence of entities which can function as truth makers for the propositions embedded in them, Cuneo is gaining ontological

commitments too cheaply. By contrast, the reason that the Moorean paradox fascinated Wittgenstein was not that it teaches us something about the 'commitments of our discourse,' but rather that it teaches us something about the nature of content itself (Wittgenstein 1972; Winch 1998). The moral of the Moorean paradox, I will suggest, is that it calls into question the metaphysical realist's picture of content, according to which content is something that is born by propositions, themselves understood as abstract entities which can be understood without reference to the attitudes of their bearer.

3. Indifferent Judges and Moore's Paradox

Before turning to broader issues concerning metaphysical realism, we must first examine the role that the Moorean paradox plays in the defence of the epistemic existence premise. To hold that ESRs exist but do not entail categorical reasons for belief means holding that the relation between facts (including probability raising facts) and beliefs is dependent on the interests of the subject in question. In Cowie's terms, facts provide us with 'non-normative,' 'institutional' reasons for forming beliefs. These are like "the reasons associated with games and sports, etiquette and the law" (2016, 121) which are supposed to apply to an individual only given their possession of a relevant interest in them.¹⁴ On this view, the existence of an ESR in favour of *p* provides grounds to believe that *p* only given that this entailment is consistent with the interests of the judger. It would be coherent for one to possess evidence for a proposition but no reason to believe it – if one were to have no "reason to engage in the business of believing (the truth) with respect to that proposition." (Cowie 2016, 121) This view is not unique to Cowie. Sharon Street gives a similar account:

Assume, for example, that all previously encountered tigers were (...) carnivorous. My claim will be that the status of this fact as a normative reason to believe that the next tiger will also be carnivorous ultimately depends on our evaluative attitudes, and in particular on the evaluative attitudes of the agent whose normative reasons are in question. (Street 2009, 218)

Cowie's rejection of the epistemic existence premise depends on the feasibility of construing epistemic norms in these terms. On this view, ESRs provide reasons for belief only given a suitable evaluative profile on the part of the believer. This profile must be independent of both the ESR and the belief, such that it would be possible for there to be an individual in relevantly similar circumstances who has access to the same ESR but who lacks any corresponding reason. The evidence that this person has for a given belief provides them with no normative guidance as to its appropriateness. If they are reflective, they may

¹⁴ Cowie implies that the members of this group are similar with respect to their normative force. Accordingly, in what follows I will take the norms inherent in games to be representative of the broader class of 'institutional norms' as a whole. Should the irrealist take issue with this, the burden is on them to specify in what the disanalogy consists.

say to themselves “*e* counts in favour of *p*, but since I don’t care about it, I have no reason to believe it.” Call this person the *indifferent judger*. The feasibility of an irrealist construal of belief depends on our being able to vindicate the possibility of such a figure.

Such vindication is immediately rendered problematic by the apparent contradiction in accepting something’s being evidence without taking it to have reason-giving force. Just as there is something amiss in accepting the truth of a claim but refusing to believe it, there seems to be something amiss in accepting evidence for a proposition without granting that one has reason to believe it. The indifferent judger thus exemplifies an analogue of the Moorean ‘absurdity’ of an individual who sincerely judges “*p* but I don’t believe it” (Moore 1942, 540–543). In the case of evidence, the analogous claim is as follows:

(E) I have evidence *e* that *p*, but I have no reason to believe it.

The challenge for the irrealist is then to defuse the challenge posed by (E) in a way which validates the possibility of the indifferent judger. At the outset it should be noted that (E) is in fact more troubling than the typical Moorean statement. The reason that Moore calls his observation a ‘paradox’ is that the state of affairs picked out by the absurd statement (viz., someone’s falsely believing that *p*) is easily conceivable. (This in turn makes it seem that the problem has something to do with the self-attribution of beliefs, or to do with semantics or pragmatics of belief assertions.) By contrast, one may doubt whether (E) expresses a possible state of affairs at all. While the relation between belief and truth is such that it is uncontroversial that a subject can believe that *p* when *p* is not the case, it is unclear whether it is possible to have evidence that *p* without reason to believe it. The irrealist’s defence must therefore involve two steps. Firstly, they must show that (E) expresses a conceivable state of affairs. Secondly, they must account for the absurdity of (E) in terms of practical rather than theoretical failure.¹⁵ Let us take these in turn.

4. The Nature of Evidence

We have seen that a defence of epistemic irrealism must begin with a demonstration that (E) refers to a possible state of affairs. Two such demonstrations may be discerned. The first proceeds by appeal to intuition – cases can be constructed where it seems natural to say that a subject has evidence for *p* but no reason to believe it. The second proceeds by appeal to the (supposed) nature of evidence.

¹⁵ See (Green and Williams 2007, ch.1) for this distinction in the context of Moorean sentences and for a useful overview of the different approaches to the problem.

Irrealists often point to so-called 'trivial truths' as evidence for the conceivability of (E).¹⁶ A trivial truth is a true judgement, belief in which has no practical relevance for a given individual (Treanor 2012). (What truths count as trivial will therefore depend on both the nature of the truth and the interests of the subject, relative to their circumstances.) Someone might have evidence for a trivial truth, but no interest in believing it. This is put forward as an instance of someone's being in the situation described by (E).

It is undeniable that truths can be more or less interesting, and as such will legitimately command an individual's attention to varying degrees. There are two lessons we can draw from this fact. The first is that world's being as it is does not in itself give the agent a reason to judge of it that it is that way. Things do not somehow 'call out' for acknowledgement.¹⁷ Secondly, inquiry is not a process of simply registering facts about the world, without discrimination due to interest or inclination. The claims we make and interrogate, and the evidence we gather and test, are affected by factors which are peculiar to us as both human beings and individuals.

But we can accept these points without conceding that the existence of trivial truths proves that (E) is conceivable. A realist ought to dig their heels in at this point and insist that if a subject has evidence for *p* then *ipso facto* they have an interest in believing it, whether they would recognize that interest or not. If interests were not objective in this sense, it is hard to see how we could avoid unpalatable conclusions such as that a child with no interest in mathematics has no reason to believe that $2+2=4$. Thus, the insights to be found in the example of trivial truths are best captured in a realistic construal of epistemic norms, as follows: although it is up to the subject whether or not to investigate whether *p* is true, if she has *e* that *p* then (it is an objective fact that) she has a reason to believe it.¹⁸

Changing tack, an irrealist may try to argue to (E) from the supposed nature of evidence. The underlying thought here is contained in the irrealist's definition of ESRs as objective (probability raising) relations that hold between states of affairs. This innocuous looking phrase implies that there is an equivalence between the statements "there is evidence *e* for *p*" and "the state of affairs picked out by *e* stands in a probability raising relation to the state of affairs picked out by *p*." In this formulation, reference to a subject is conspicuous by its absence. Evidential relations are subsumed under a quasi-causal model, as

¹⁶ This is somewhat of a misnomer since this is not the logician's sense of 'trivial'. Rather, truths are 'trivial' in the relevant sense if they are pieces of trivia, i.e. uninformative or uninteresting.

¹⁷ This is what Street calls the denial of 'Cliffordian normativity' (2009, 218-221).

¹⁸ An analogy can be drawn with the moral case, as follows. It is up to the individual whether or not to incur certain special duties, e.g. whether or not to make a promise to someone else. However, once the duty is incurred, it then binds on that individual categorically, without reference to their desires.

relations between states of affairs. From here, what it is to possess a piece of evidence can be defined as follows: A has *e* for *p* iff A stands in a suitable relation of epistemic access to *e* (for instance, its being the case that A would be aware of *e* if she attended to it). On this picture, in judging that there is *e* for *p* we are simply registering a connection that holds between two subject-independent features of the world (two facts). It seems queer to think that such a relation should (uniquely amongst natural relations) be automatically reason-conferring. This therefore licenses the third person analogue of (E):

(E*) A has *e* for *p* and A has no reason to believe that *p*.

But if (E*) describes a coherent state of affairs, then the state of affairs in (E) must be coherent also, as A in (E*) and the first personal pronoun in (E) are co-referring. Thus is the analogy with the Moorean paradox to be secured.

One way to push back against this argument would be to argue that epistemic access relations bring with them categorical norms; in other words, to hold that registering probabilistic connections between facts involves commitment to acting in certain ways.¹⁹ But rather than pursue this possibility, I would rather call into question the account of ESRs as objective relations holding between states of affairs. This account takes for granted that propositions are related to the world in such a way that things, by virtue of being as they are, call for a certain determinate set of judgements. Following Putnam, let's call this view *metaphysical realism*. Without this assumption we could not speak of 'the state of affairs picked out by *p*' as standing in an objective relation with something else. On the metaphysical realist account, truth is a matter of correspondence between two states of affairs; a representing fact (the judgement) and a fact represented (a state of the world).²⁰

The fine details of such a picture may vary, but the broad strokes will not, and it is the broad strokes with which I will take exception. The following illustration will therefore suffice for my purposes. Take two states of affairs. To say that these stand in a relation of probabilification is to hold that if the former obtains then the chance of the latter obtaining is modified to some degree *n*, where *n* ranges between 0 (impossible) to 1 (certain). A limiting case of probabilification is entailment, where one state of affairs makes another metaphysically certain. Along with relations of probabilification there are

¹⁹ See, for instance, McDowell's contention that judgements of matter of fact bring in reasons for action in the activity of conceptualization rather than in their content (1979).

²⁰ For a contemporary endorsement of this sort of 'correspondence view' (albeit by moral realists) see (Cuneo and Shafer-Landau 2014). Both parties to the contemporary realism/anti-realism debate share a (usually implicit) commitment to both the correspondence view and the metaphysical realism that underwrites it. The implausibility of fitting categorical norms into the world as conceived in metaphysical realist terms leads anti-realists to reject the possibility of categorical normativity. My aim is to show how objectivity may be rehabilitated not through the vindication of a correspondence view but rather through the rejection of the metaphysical realism that makes it seem necessary.

relations of instantiation. According to the picture we are sketching, the instantiation relation is in fact a special case of the entailment relation. Things being as they are instantiates some propositions (i.e. some ways that things can be described as being) and not others. If instantiation relations can be understood as objective relations holding between judgements and states of affairs then we may say that, for a given state of affairs, there is a given class of judgements which are true of it. A complete description of a state of affairs consists of the conjunction of all general descriptions which it instantiates. Each of these judgements will capture some aspect of the state of affairs, potentially relevant to the truth of some judgement with respect to those further states of affairs to which this state stands in a relation of probabilification.²¹ In this way we can move from objective relations between states of affairs to objective relations between a given state of affairs and a given class of judgements, and so in turn to a characterisation of evidential support relations in non-subject-involving terms.

On this picture, the obtaining of the fact *all previously encountered tigers were carnivorous* makes it objectively more likely that the fact *this tiger is carnivorous* will obtain.²² In turn, the obtaining of this latter fact objectively warrants the judgement that "this tiger is carnivorous." Since both connections are objective, there is thus an objective connection holding between the fact that *all previously encountered tigers were carnivorous* and the judgement "this tiger is carnivorous." In this way, the appropriateness of a given factual judgement is determined simply by how the world is, and the role of the judge is confined to making themselves receptive to relations which hold between facts.

With this in place epistemic subjectivism heaves into view, exploiting metaphysical realism's distinction between truth and the norms of truth-tracking. If truth is defined in terms of a content-bearing state's standing in an appropriate relation with the state of the world that answers it, then we are free to see the norms of truth-tracking as garnish; desirable but dependent on the subject's interests.

However, and crucially, this argument for epistemic irrealism turns on an acceptance of metaphysical realism. In the absence of such a commitment there would be no way to move from objective relations between features of the world to objective relations between the world and propositions concerning it. In other words, we will not be warranted in seeing the instantiation relation as a limiting case of entailment. If we reject this conception then the irrealist's conception of evidence falls by the wayside. The conceivability of (E) thus cannot be

²¹ Generalising, we can represent the class of true judgements which this state of affairs instantiates as standing in objective probabilification relations with the total set of all the sets of judgements instantiated by all and only the states of affairs probabilistically related with this one.

²² I am bracketing here substantial complications to do with the status of facts involving quantification, indexicalisation and temporal reference.

demonstrated in advance of settling broader controversies surrounding the nature of truth and evidence. So, appeal to (E) in defence of epistemic subjectivism can only beg the question.

5. Practical Versus Theoretical

So far I have shown that the conceivability of (E) – and thus the fate of epistemic subjectivism – depends on the truth of metaphysical realism. I now turn to attempts by the epistemic subjectivist to explain the problematic nature of (E) in terms of practical inconsistency. The most promising of these accounts invokes normative requirements that are built into the concept of believing but that do not bear on belief's being a representational state.²³ On this view, it is a criterion of believing that *p* that one takes evidence that *p* to give one reason to believe it. The introduction and elimination rules for belief simply do not permit the acknowledgement of evidence without its bearing on belief, any more than the rules of chess permit castling through check. But, the irrealist insists, this framework in turn binds on us only because of our interest in it. The requirement that beliefs track truth is, in Street's terminology, merely a 'conceptual' rather than an 'objective' requirement (Street 2009, 243).

What it means to call belief an 'institution' or a 'game' is intimately bound up with this contrast between conceptual and objective requirements. It depends on our being able to see the requirement in question as in some sense non-compulsory; such that an individual, in the same circumstances, might without error comport themselves according to a different rule. One way of expressing this point is to imagine an alternative conceptual scheme with different conceptual requirements, and to demonstrate that the (hypothetical) choice between these two schemes can only be made on pragmatic grounds.

In the case of a game such as chess, that would mean having a variant on chess – call it chess* – in which (say) castling through check is permitted. In the case of belief it would be for there to be a rival concept of belief – call it belief* – on which evidence in favour of *p* does not provide a subject with a reason to believe* that *p*. If it can be shown that the only grounds for preferring chess over chess* are pragmatic ones, it would follow that there is no cognitive deficiency in preferring the chess* concept. And so, *mutatis mutandis*, for belief. In that case, our adherence to belief rather than belief* could only be justified on practical grounds.²⁴

However, chess* counts as a game only if it can be played, in other words, only if it can be substituted for chess without undermining the status of

²³ Other approaches include the claim that (E) leads to a kind of pragmatic contradiction arising from the conversational implicatures of assertion. Further elaboration and criticism of such an account can be found in (Heal 1994).

²⁴ For example, in terms of the greater evolutionary success of beings whose beliefs track evidence (Street 2015, 240).

participants as players of a game. Analogously, belief* counts as a cognitive state only if it can be substituted for belief in an individual's conceptual repertoire without undermining their status as a thinker. Thus, the plausibility of this irrealist argument turns on the tenability of the belief* concept. Beliefs* must represent states of affairs without encoding any norms governing relations between belief and evidence. The belief* concept is plausible only if there can be a representational state the content of which is settled without reference to the attitudes of the judging subject. And so, the belief* concept and metaphysical realism stand and fall together.

In this way, both parts of the defence of epistemic subjectivism depend on the truth of metaphysical realism. This should be unsurprising, for in order to deny that the failure in (E) is 'theoretical,' one must either deny the existence of evidential relations entirely, or else slough them off into the realm of the objectual. In this way, the doctrine of (E) as a form of practical failure is simply a variant on the thesis that epistemic access relations are subjective relations, logically derivative on the prior, non-subject-involving relation of something's being evidence for something else.

6. Metaphysical Realism and Epistemic Irrealism

We are now in a position to see how epistemic irrealism might be undermined. Here is the argument in a nutshell. If the content of a belief always depends implicitly on the attitudes of the believer (specifically, on what they are willing to countenance as evidence in favour of their belief), then to hold that A has *e* for *p* is to make a claim which is implicitly subject involving: it depends for its truth on A's acceptance of a particular understanding of what counts as *p*. Clarifying the meaning of an utterance means determining more precisely range of circumstances under which it would be true. This fixes in turn what can count as evidence for the proposition. Accordingly, to take *e* to bear on some claim is *inter alia* to endorse one understanding of *p* out of many possible. The norms of evidence are therefore norms of representation. And so the absurdity in (E) stems from its being self-contradictory; to attempt to have a representational state without the norms of representation is to try both to have one's cake and eat it.

This view of content involves the rejection of metaphysical realism, and is supported by arguments that show the implausibility of the metaphysical realist account of content. Recall that the metaphysical realist holds that the content of content-bearing states can be fixed independently of the attitudes of the judger.²⁵ When a state's content gives it truth conditions, then its being true (or false) is determined entirely by the ordered pair of the properties of the content-bearing state as related to the properties of some state of the world.

²⁵ A natural way to imagine this is to think of belief as a sub-personal state of an agent. But we do not have to make this supposition in order to accept the general point.

Bearing this in mind, we can construct an argument against the feasibility of metaphysical realism as follows. Take the judgement that all tigers are carnivorous. Whether it is true depends, *inter alia*, on what is meant by 'tigers.' Suppose that all large cats on earth are carnivorous, but that on Alpha Centauri there lives a species of herbivorous large cats, genetically similar to tigers but evolutionarily unrelated to them. Imagine now that upon returning from a jungle expedition, Abe declares "all tigers are carnivorous." Did she speak truly? This depends on whether the extension of her expression 'tigers' includes all large cats genetically similar to earth-tigers, or whether it includes earth-bound species only.²⁶

Determining whether the large cats of Alpha Centauri count as tigers is part of the process of clarifying the content of Abe's belief. How are such matters to be settled? We can imagine two different concepts of 'tiger,' one of which includes space cats within its extension and the other of which excludes them. Before she had encountered Alpha Centauric creatures, Abe could not have had those particular individuals in mind when she made her judgement (Travis 2006, 130-135). She could thus have intended neither to include nor to exclude them from consideration. Hence, no property of Abe at the time of her utterance settles which concept is in play in her judgement.

We can use the very same form of words to express different thoughts on different occasions. In certain circumstances, the thought that Abe expresses with the statement "all tigers are carnivorous" will be such as to turn on the dietary habits of creatures from Alpha Centauri (if, for instance, she forms an extra-terrestrial exploratory team which is then confronted by a hungry looking space-cat.). On other occasions, the thought may turn only on the properties of animals that are closer to home. Her intentions in making the judgement are relevant to determining which thought she expresses, hence whether or not her judgement is true. In this way for any given belief the content of the belief depends in part on the attitudes of the believer, and in particular on what they are willing to countenance, on some particular occasion, as things being as they took them to be.²⁷

²⁶ For an elaboration of this point see (Putnam 1975; Travis 2011, 130-143). Although the case is outlandish, the general insight is not. As Travis puts it: "For any way there *is* for things to be, things being as they are *may*, in point of grammar, count as that way on some understandings of so being and not that way on others." (2011, 135).

²⁷ See (Travis 2006, 26ff) for an elaboration of this point, in particular: "'is blue' does not, in naming what it does (in speaking of (something's) being coloured blue), acquire the function of being true of things. That is *not* its function. The function it acquires is, rather, that of speaking, on an occasion, of something's being blue when used, on that occasion, as naming what it does. It may do that while functioning as a move in any of indefinitely many language games. Which is, here, just an imagistic way of saying that it may do that while subject to any of indefinitely many mutually incompatible conditions on the correctness (truth) of what is thereby said." (2006, 28-9).

But although Abe's intentions are relevant, they are not decisive. In responses to novel cases we are constrained by our answerability to the community (Travis 2011, 10-12). Not everything can count as a reasonable interpretation of what it is to be a tiger; not everything can count as a reasonable response to a novel case. (Compare: "by 'tiger' I mean 'all felines except those that don't eat meat'.") In judging that p one incurs normative commitments; namely answerability to shared standards for determining, on a given occasion, whether things being as they are counts as an instance of things being that p or not.

We have here in highly compressed form a familiar argument for the irreducible normativity of content. The normativity comes from the fact that content is made possible only by virtue of our standing as answerable to the indefinite community of thinkers, as that which ensures that some understandings of the rules by which we proceed are ruled in and others ruled out. The irreducibility comes from the fact that any specification of the extension would be open to a further question of interpretation. If these standards cannot be given a reductive specification, there is no possibility of eliminating normative commitments in the determination of the content, for some given occasion, of the speaker's judgement.²⁸

This is a point about the specification of content in the first instance, and not about the relationship between a given belief and evidence that bears on it. But the evidence that counts in favour of a given belief is dependent on the belief's content, in the following way: to specify what counts as thing's being in accordance with a belief is to delimit what counts as evidence in favour of the belief. If, for instance, we specify that only creatures with appropriate evolutionary links to our paradigm large cats count as tigers, then facts about evolution become evidence that bears on the truth or falsity of the claim that all tigers are carnivorous. *Mutatis mutandis*, if we treat facts about evolution as relevant to settling of an object whether or not it is a tiger, then our concept of tiger is a biological category, such as to rule out of court a stuffed animal of the right appearance counting as an instance of the relevant kind. Thus, to present e as evidence for p is to hold that p ought to be understood in one way rather than another.²⁹

²⁸ One argument against the possibility of such a reduction runs as follows. A reductive specification of normative standards would lead to an infinite regress: take any specification of a rule for determining whether or not p holds. We may then ask "what counts as things being in accordance with that rule?" Any further specification will simply invite the same question (Wittgenstein 2009, §185ff.).

²⁹ This view ought not to be confused with the doctrine of evidentialism, which in its strongest form holds that to believe that p entails judging oneself to have evidence adequate to the truth that p and *vice versa*. (Adler 2002: 32) On this view one's evidence exhaustively determines the content of one's beliefs. But this is surely false. It seems clear that one can believe that p , without having any evidence for the truth of p and even without having any very clear notion of what could count as evidence for or against p . This is one of the chief insights of Putnamian

How does this relate to the putative normative requirement that if one accepts that one has evidence for a proposition, then one must take oneself to have a reason to believe it?³⁰ To present something as evidence is to say that things are a certain way, and that in being that way, they are another way too (if there is a paw print in the mud, then there is a tiger nearby.). Thus in presenting evidence for some claim we fix an understanding of it by endorsing it. This seems to be part and parcel of how evidence works, and explains how it is that we can determine not only the truth of a belief but also its content by interrogating what is to be counted as evidence in favour of it. If this is correct then we cannot hold onto the notion of evidence while suspending commitments to beliefs across the board.

7. Belief and Institutional Norms

When Abe encounters Alpha Centauric space cats she learns something new about the world. Suppose that experts agree that these creatures are indeed tigers. Then her knowledge is a discovery about the nature of tigers; namely, that some of them are herbivores.³¹ The possibility of this kind of discovery depends on the determination of the extension of our concepts being to some extent out of our own control. This in turn helps to explain why we cannot define the norms of belief as we do institutional norms such as games. In the latter case, the question may, in a given circumstance, arise as to whether a certain state of affairs counts as a violation of a given rule. In such cases, a judgement call will be required. In that respect institutional norms also depend on the attitudes of the participants; rules need to be applied (Wittgenstein 2009, §185ff.). But although there is similarity between the cases in that respect, there are also crucial

semantic externalism. If meaning were fixed by psychological states (whether of the subject or of the community as a whole), then firstly communication between people with different standards of verification would be impossible, and secondly genuine discoveries about the world would be impossible. See (Diamond 1999) for exposition of this strand in Putnam's work. We could not discover, say, that water is in fact H₂O since that identity, by changing the methods of verification, would change the meaning. By contrast to Adler, all I am committed to here is the weaker claim that to fix the content of one's belief is to specify what counts as evidence for it, and *vice versa*. This is enough to ensure that possession of a belief entails commitment to a norm of evidence, presuming that there is an internal connection between 'having a belief' and being prepared to specify its content.

³⁰ Using 'normative requirement' in John Broome's sense. A normative requirement differs from a reason in the following sense: if *e* is not in fact evidence for *p* then one has no reason to believe that *p*, but if one takes something to be evidence for *p* and yet fails to take oneself to have a corresponding reason to believe that *p*, one violates a normative requirement and so is "not as one ought to be." (Broome 1999)

³¹ Compare; learning that water is H₂O was a discovery about the nature of water. As a result of this discovery our use of the term 'water' changed; we can now distinguish between water and superficially similar liquids (like Twin Earth's XYZ) (Putnam 1975).

differences. One difference which is particularly important is that in the case of a game settlement of novel cases can only proceed through fiat.

Abe thinks that Alpha Centauric cats are tigers; Mie disagrees. There are indefinitely further considerations that each can appeal to in settling the issue; various branches of science, common knowledge, the purpose for which the judgement is being made, and so on. Any particular interpretation of the content of a belief will have knock on effects on other beliefs. Cognitive states ramify indefinitely. Suppose Alpha Centauri large cats have blue coats. Abe is therefore committed (whether consciously or not) to the proposition that some tigers have blue coats. Thus to the proposition that some animals do. And so on. This prevents Abe's reaction to a novel case from being arbitrary. When she decides whether to endorse or abandon her belief, she must weigh up the relative costs and benefits of each option, comparing the losses in terms of disruption to the established order, to benefits in terms of a new capacity for ordering or coping with novel or recalcitrant experiences.³²

By contrast, suppose that in a game a controversy arises as to the interpretation of a certain rule. Abe thinks that the ball is out of play if caught by a player in the air over the line of touch. Mie disagrees, holding that the ball is only out of play if that player makes contact with the ground out of the field of play. Abe may appeal to any number of considerations in favour of her proposal, for instance that a game played by her rules would be more entertaining. But the only considerations that Mie must accept as relevant to the issue are those provided for by the nature of the game itself, and *a fortiori* these do not settle the issue. If Mie chooses to take Abe's considerations as relevant to the issue at hand, that is a matter of her sentiments – of her sharing with Abe a conception of the importance of the game's being entertaining, or elegant, or what have you. Since neither alternative is forced on either Abe or Mie, neither is vulnerable to the accusation of perversity.³³

This divergence has knock on effects on the status of disagreements between players of games and between believers, respectively. A disagreement constitutes a problem in a game only if it is a practical impediment to playing

³² There are of course a number of different ways of unpacking this holistic requirement. A *locus classicus* is (Quine 1951). But, as Travis (2011, 101) warns, we must be careful not to hold that any belief may be accepted given radical enough changes to surrounding assumptions and come what may. The upshot of this would be to deprive ourselves of the notion of a content to belief (of its being *about* something) entirely. Thus in some (perhaps most) cases, there will be nothing else for an individual to think other than that things are as they appear on face value to be. For attempts to reconcile this role for experience with holism see (Travis 2006, 30; McDowell 1996).

³³ Or, rather, the sense in which one may be accused of 'perversity' changes. One wants to say that in the belief case there is an impersonal necessity which resolves conflicts between disputants. But in the case of a game, human nature may make certain arrangements and not others compelling, and this too may have the force of an 'impersonal necessity' (if, for instance, beings such as us cannot help but find games of this kind compelling).

and, accordingly, the resolution of the issue can only be a pragmatic matter. Suppose Abe and Mie cannot agree about the interpretation of the 'out of play' rule in their game. The resolution of the issue would involve a further specification of the rules to cover problematic cases, for instance: "the ball is deemed out of play only when it touches the ground outside of the field of play." If Abe does not agree to the new rule, then she can go off and form a new group, playing with her own house rules. Abe and Mie would then no longer be disagreeing; they would simply be playing different games.

On the other hand, suppose that Abe and Mie cannot agree over the extension of the concept 'tiger'; over whether that blue coated space-cat counts as a tiger or not, and thus (generalising) over the truth of such claims as that "some tigers have blue coats." There is no analogue to the concept of 'house rules' in this case. Unlike in the case of the game, the possibility of controversy arising between two competent subjects does not call into question it being the case that they are speaking of the same thing.³⁴ Thus, rather than settling controversies before conversation can proceed, points of controversy become points of conversation. Even if the discussants conclude that they mean two separate things by 'tiger,' and resolve by disambiguating, neither goes off to "play their own game." Rather, both senses are incorporated into their common conceptual repertoire.³⁵

These disanalogies cast doubt on the adequacy of an account of the norms of belief as merely 'conceptual' requirements. The deployment and refinement of a game's concepts can only ultimately be justified by reference to our own amusement. In the case of beliefs about the world, the deployment and refinement of our concepts is geared towards the growth of a subject's understanding. And this difference, in turn, explains the different status enjoyed by normative requirements in the two contexts.³⁶

This concludes my defence of the epistemic existence premise. I have tried to show that the irrealist is tacitly endorsing the following principle: either the world, solely in virtue of being as it is, gives us reasons to make certain judgements and not others, or else the norms governing the making of judgements are merely subjective. I have argued that the correct response to this challenge is not to try to rehabilitate the notion of an 'objective forcing' (understood in the irrealist's terms) but rather to deny the dichotomy. Then, we may say that although the world alone does not make any particular judgement

³⁴ This raises difficult issues concerning the limits of the analogy between languages and games. See (Wittgenstein 2009, §1ff.; Rhees 1959). A relevant point here is that we can speak of translation between languages whereas we do not talk about translation between games.

³⁵ This is part of the process by which we grow in understanding through conversation; a phenomenon which has no analogue in the case of a game (Rhees 1959).

³⁶ Thus games are constrained by the laws of nature, but are not answerable to them. A legitimate move in a computer game may involve a character within the game violating the laws of physics.

obligatory, it doesn't follow that the choice of whether or not to hold a certain belief on some given occasion is dependent on the subject's interests.

On the rival picture that I have (very roughly) sketched, things, in being as they are, give us neither a categorical nor a hypothetical reason to form any particular belief, because any way for things to be is, in principle, consistent with a range of different and mutually exclusive beliefs.³⁷ Nevertheless, there is a relation between the world's being as it is and our beliefs, as follows: when the question of whether or not to endorse a particular belief is opened, or when the world, in being as it is, is taken as evidence for a certain belief, then whether to form or maintain that belief is not dependent on the subject's interests. Things being as they are, given the background on which the judgement occurs, delimits the range of beliefs appropriate to it. If we are lucky it may even determine, exclusively, one belief or set of beliefs as the only ones that can reasonably be held. As Putnam stresses, this determination depends on both the discretion of the subject and their background of shared commitments. But to depend on the discretion of the subject is not the same as being relative to the subject's interests, for the scope of the subject's discretion is constrained by the overall aim of understanding.³⁸

8. Reduction and Grounding

This brings us, finally, to the argument in favour of the parity premise. Here, we can take a shorter line. Recall that the 'Master Argument' takes the form of a dilemma. Its two horns represent two kinds of moral irrealist. Those who choose the first horn deny the existence of categorically normative reasons across the board. Those who choose the second horn are committed to the existence of some categorically normative reasons, but refuse to include moral reasons amongst that class. We have already seen that the former position is untenable; the denial of the existence of categorical reasons for belief is self-contradictory.

That leaves, then, the latter position. This view is concessive in that it accepts that categoricity *per se* is not a bar to the objectivity of a given judgement, so long as that judgement possesses some further feature to render it unproblematic. In this case, the further feature, apparently possessed by

³⁷ Nor is there a reason to form some belief (understood non-specifically), as the case of trivial truths shows us.

³⁸ Of course, a comprehensive defence of the Putnam-Travis view would have to do more. In particular it would have to address the following two concerns. Firstly, cases can be constructed where it seems that an individual has a belief even though nothing would count as evidence for or against it (belief in God may be an example of this.). Secondly, it may be urged that animals can form beliefs about their environment even though they lack the ability to reflect on the evidence base of their beliefs or on its content. If the Putnam-Travis view is correct an explanation must be given of these cases which is consistent with the claim that there is an internal relation between a state's being representational and the judger's commitment to evidential norms.

epistemic but not moral judgements is the possession of a suitable 'reductive-base'; epistemic reasons can be reduced to some collection of non-normative facts or properties.³⁹

I suspect that only a tacit acceptance of metaphysical realism could make a reduction of epistemic norms seem feasible. Reasons are relational properties, one term of which is a (judging or acting) subject. Probabilification relations that hold between two facts or states of affairs make no reference to subjects. For this reason, any account of the former in terms of the latter would seem of necessity to be incomplete.⁴⁰ The only grounds for holding that reasons for belief can be defined in terms of probabilities that hold between states of affairs, it seems, is a prior conviction that reality can be exhaustively characterised in terms of a collection of mere happenings. But, as Putnam himself has long urged, it is hard to see how that picture could make room for the notions of judgement or truth at all (Putnam 1981, ch.2).

There is a second problem for this form of concessive irrealism, related to the first. As Heathwood notes, if we were to pull off a reduction of epistemic reasons to collections of non-normative facts, the upshot would seem to be scepticism about epistemic normativity:

If epistemic normativity is characterizable in purely descriptive terms, this suggests that, in some sense, epistemic normativity is not *real* normativity. If some belief is unreasonable, we do not *really have* to stop holding it. True, it would be epistemically irrational not to, but this is to say no more than that the belief is not likely to be true; it is not to say that we *really must* stop believing it. (Heathwood 2009, 96)⁴¹

In this way, the postulation of a reduction-base for epistemic norms sets us back into either nihilism or subjectivism, both of which are non-starters. An awareness of this danger may explain the ambivalence that error theorists sometimes display towards the prospect of providing a 'reductive-base' for epistemic norms. It is sometimes suggested that the relevant relation may be not 'reduction' but 'grounding'. But an explanation is then owed of the nature of this 'grounding' and its relevance to the metaphysical premise of the error theory. For, many, if not most, moral realists will accept that moral norms are (in some sense) grounded in natural facts. For instance, most moral realists accept at least a weak form of supervenience, according to which there can be no change in the

³⁹ One such reduction, suggested by Chris Heathwood and endorsed by Cowie, is an account of epistemic reasons as facts about likelihoods, themselves characterisable in non-normative terms (Heathwood 2009, 92).

⁴⁰ This point is obscured by Cowie's use of the concepts of 'evidence' and 'hypothesis,' both of which make implicit reference to a subject, and are, for that reason, not fit candidates to feature in a naturalistic account of probabilities.

⁴¹ Note, however, that Heathwood himself does not find a denial of the reality of epistemic norms troubling.

distribution of moral properties without a corresponding change in the distribution of natural properties.⁴²

In conclusion, in the absence of a satisfactory account of the 'reductive base' of epistemic norms, the concessive error theoretic position fails to mark out any relevant distinction between epistemic and moral reasons. It therefore fails to provide grounds for rejecting the existence of the latter. In other words, the parity premise is secure.

9. Conclusion

For the reasons canvassed above, Cowie's 'Master Argument' against the CG strategy is unsuccessful. But there is a broader moral here. The point of the CG defence is to bring our standards for objectivity back to earth, by showing to be overly demanding the metaphysical strictures upon which the irrealist relies. Cowie thinks that the irrealist can avoid this charge, either by 'biting the bullet' and denying the existence of categorical reasons for belief, or else by finding some further, metaphysically reassuring feature unique to epistemic reasons. I have shown that both of these responses rely on the same questionable metaphysical picture of what it takes for a judgement to be objectively true. What is really doing the work in the CG argument is thus an underlying debate about the relationship between the world and our judgements about it. Once this is accepted, we will see that there is no need to head off into the hills in pursuit of some further metaphysically reassuring feature possessed by moral reasons. No metaphysical reassurance will be forthcoming, but none is required.

What, then, ought we say about the relation between ESRs and categorical epistemic reasons – are they identical, or not? The question conflates identity as understood at the level of types and at the level of tokens. There is no type-identity between ESRs and categorical reasons for belief, and thus no possibility of reducing one category to the other. However, there may be an identity that holds between token instances of ESRs and categorical reasons, as follows: if *e* provides evidential support for *p*, then for a suitably placed subject, *that very fact* is a categorical reason in favour of their believing that *p*. But this identity holds only given an understanding of *p*, which is an occasion-sensitive matter. In this way the subject, in meaning their words, transforms facts into reasons.

References

- Adler, Jonathan. 2002. *Belief's Own Ethics*. Cambridge: MIT Press.
 Blackburn, Simon. 1993. *Essays in Quasi-Realism*. New York: Oxford University Press.
 Broome, John. 1999. "Normative Requirements." *Ratio* 12 (4): 398–419.

⁴² This is a very rough statement of supervenience; for a more careful discussion see (Shafer Landau 2003, ch.4).

Michael Campbell

- Cowie, Christopher. 2014. "Why Companions in Guilt Arguments Won't Work." *The Philosophical Quarterly* 64 (256): 407-422.
- . 2016. "Good News for Moral Error Theorists: A Master Argument against Companions in Guilt Strategies." *Australasian Journal of Philosophy* 94 (1): 115-130.
- Cuneo, Terence & Shafer-Landau Russ. 2014. "The Moral Fixed Points: New Directions for Moral Nonnaturalism." *Philosophical Studies* 171 (3): 399-443.
- Cuneo, Terence. 2001. "Are Moral Qualities Response-dependent?" *Noûs* 35 (4): 569-591.
- . 2007. *The Normative Web: An Argument for Moral Realism*. Oxford: Oxford University Press.
- Diamond, Cora. 1999. "How Old Are These Bones? Putnam, Wittgenstein and Verification." *Aristotelian Society Supplementary Volume* 73 (1): 99-134.
- Dreier, James. 2004. "Meta-ethics and the Problem of Creeping Minimalism." *Philosophical Perspectives* 18 (1): 23-44.
- Evers, Daan. 2015. "Street on Evolution and the Normativity of Epistemic Reasons." *Synthese* 192 (11): 3663-3676.
- Foot, Philippa. 1972. "Morality as a System of Hypothetical Imperatives." *Philosophical Review* 81 (3): 305-316.
- Green, Mitchell S. & John, Williams, eds. 2007. *Moore's Paradox: New Essays on Belief, Rationality, and the First Person*. Oxford: Oxford University Press.
- Heal, Jane. 1994. "Moore's Paradox: A Wittgensteinian Approach." *Mind* 103 (409): 5-24.
- Heathwood, Chris. 2009. "Moral and Epistemic Open-question Arguments." *Philosophical Books* 50 (2): 83-98.
- Joyce, Richard. 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.
- Lillehammer, Hallvard. 2007. *Companions in Guilt: Arguments for Ethical Objectivity*. Basingstoke: Palgrave Macmillan.
- Littlejohn, Clayton. 2016. "Do Reasons and Evidence Share the Same Residence?" *Philosophy and Phenomenological Research* 93 (3): 720- 727.
- McDowell, John. 1979. "Virtue and Reason." *The Monist* 62, 331-350; reprinted in his (1998).
- . 1985. "Values and Secondary Qualities." In *Morality and Objectivity*, edited by Ted Honderich, 110-129. London: Routledge.
- . 1996. "Two Sorts of Naturalism." In *Virtues and Reasons: Philippa Foot and Moral Theory*, edited by Rosalind Hursthouse, Gavin Lawrence, and Warren Quinn, 149-79. Oxford: Clarendon Press. Reprinted in *Mind, Value and Reality*.
- . 1998. *Mind, Value and Reality*. Cambridge, Mass.: Harvard University Press.

- Moore, G. E. 1942. "Reply to My Critics." In *The Philosophy of G. E. Moore*, edited by Paul A. Schilpp. La Salle, Illinois: Open Court.
- Putnam, Hilary. 1975. "The Meaning of 'Meaning'." *Minnesota Studies in the Philosophy of Science* 7: 131-193.
- . 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.
- . 1982. "Beyond the Fact – Value Dichotomy." *Critica* 14 (41): 3-12, reprinted in *Realism with a Human Face*.
- . 1990. *Realism with a Human Face*. Cambridge MA: Harvard University Press.
- Quine, Willard V. O. 1951. "Two Dogmas of Empiricism." *Philosophical Review* 60 (1): 20-43.
- Rhees, Rush. 1959. "Wittgenstein's Builders." *Proceedings of the Aristotelian Society* 60: 171-186.
- Schroeder, Mark. 2008. "What Is the Frege-Geach Problem?" *Philosophy Compass* 3 (4): 703-720.
- Shafer-Landau, Russ. 2003. *Moral Realism: A Defence*. Oxford: Oxford University Press.
- Street, Sharon. 2009. "Evolution and the Normativity of Epistemic Reasons." *Canadian Journal of Philosophy* 39 (sup1): 213-248.
- Streumer, Bart. 2013. "Can We Believe the Error Theory?" *Journal of Philosophy* 110 (4): 194-212.
- Travis, Charles. 2006. *Thought's Footing: A Theme in Wittgenstein's Philosophical Investigations*. Oxford: Oxford University Press.
- . 2011. *Objectivity and the Parochial*. Oxford: Oxford University Press.
- Treanor, Nick. 2012. "Trivial Truths and the Aim of Inquiry." *Philosophy and Phenomenological Research* 89 (3): 552-559.
- Winch, Peter. 1998. "Judgement: Propositions and Practices." *Philosophical Investigations* 21 (3): 189-202.
- Wittgenstein, Ludwig. 1972. *On Certainty*, edited by G.E.M. Anscombe and P.T Geach. New York: Harper & Row.
- . 2009. *Philosophical Investigations*. 4th edition, edited by Peter Hacker and Joachim Schulte. London: Wiley-Blackwell.