

Ethical Evidence

Penultimate draft

Forthcoming in *Synthese*

Steven Diggin, The University of British Columbia

steven.diggin@gmail.com

12,421 words.

This paper argues that ethical propositions can legitimately be used as evidence for and against empirical conclusions. Specifically, I argue that this thesis is entailed by several uncontroversial assumptions about ethical metaphysics and epistemology. I also outline several examples of ethical-to-empirical inferences where it is extremely plausible that one can rationally rely upon their ethical evidence in order to gain a justified belief in an empirical conclusion. The main upshot is that ethical propositions can, under perfectly standard conditions, play both direct and indirect evidential roles in (social) scientific inquiry.

Keywords: values in science; metaethics; moral epistemology; evidence; philosophy of social science.

Recently, I was looking after my friend's children, Spencer and Emily. I was sitting in the living room, when Emily ran out from the kitchen, where she and Spencer were supposed to be sharing a piece of cake. 'Spencer didn't cut the cake fairly!', Emily cried. I believed her, and so I inferred from this fact (the fact that Spencer didn't cut the cake fairly) to the conclusion that he didn't split the piece of cake evenly, in order to take the bigger piece for himself. I went into the kitchen to talk to both of the children – Spencer especially – about the importance of fairness and equality (without much success).

Later, I noticed that I had made an interesting inference to reach the (true) conclusion that Spencer didn't split the cake evenly. I inferred from an ethical premise, namely, that

Spencer did something unfair, to reach an empirical conclusion about the geometry of the cut cake. Although not quite deductive, this had all the hallmarks of an epistemically virtuous inference. In other words, the ethical proposition, *that Spencer cut the cake unfairly*, seemed to be the good evidence on the basis of which I believed the empirical conclusion that Spencer cut the cake unevenly.

While exploring the more general possibility that ethical¹ propositions² can, under standard conditions, be used as evidence for and against empirical conclusions, I was surprised to see (both in the philosophical literature and also especially in conversation with colleagues) that many find this idea quite absurd.³

¹ By 'ethical propositions', I mean propositions which feature evaluative or normative words like, 'good', 'bad', 'right', 'wrong', 'just', 'fair' and so on. More stringently (and in order to avoid complications involving disjunctive propositions), we can define an ethical proposition as one whose truth *essentially* depends on whether an ethical property is instantiated (or some anti-realist-friendly equivalent of this). 'Empirical propositions' are propositions whose truth essentially depends on the instantiation of the kind of property which would be described by our best scientific theories (although in this paper I will only be concerned with empirical propositions which purport to describe the world at the level of description of the social sciences). Propositions which feature "thick" predicates (e.g., 'selfish', 'admirable') plausibly fall somewhere between ethical and empirical propositions as I have defined them. If the argument in this paper succeeds in showing that ethical propositions can be used as evidence for and against empirical propositions, then it shows *a fortiori* that "thick" propositions can also be used as evidence in this way.

² Throughout this paper, I assume that bodies of evidence are composed of *propositions*, so it would be the proposition *that Spencer cut the cake unfairly* which is my evidence for the conclusion *that Spencer cut the cake unevenly*, rather than my mere *belief* about this proposition, or *experience* of the unfairness, etc. Not much turns on this assumption. The discussion below could be easily adapted to show that ethical beliefs (rather than ethical propositions) can be legitimately used as evidence for empirical conclusions. The matter is a bit trickier if we think of evidence as constituted only by experiences or immediate sense-data, since it is not obvious that we have distinctively ethical experiences or sense-data. In this case, however, it would be sufficient for the discussion in this paper to establish that ethical propositions can play the same pseudo-evidential role which empirical propositions standardly play, e.g., when these empirical propositions are learned via testimony. To state the overall point of this paper as abstractly as possible, the idea is that – whatever account of evidence one prefers – there are no grounds for treating ethical propositions or beliefs as failing to play whatever evidential or quasi-evidential role which one ascribes to empirical propositions or beliefs within this account of evidence.

³ The main location where there has been discussion (or rather, a conspicuous lack of discussion) about the idea that ethical propositions can count as evidence for empirical conclusions is the large philosophical literature on the role of value-judgements in scientific inquiry. In this debate, the distinction is commonly drawn between 'direct' and 'indirect' roles (Longino 1990, Douglas 2000:564) which value-judgements – i.e., ethical beliefs – can play in scientific research, and it is usually assumed without argument (or at most a gesture towards a worry about wishful thinking, see §1 below) that value-judgements cannot play a direct evidential role in influencing scientific research. However, there have been some notable

From now on, I refer to the principle that ethical propositions can, under perfectly standard conditions (i.e., when they meet the criteria in our best general theories of evidential support), be legitimately relied upon as evidence for and against empirical conclusions as the *Ethical Evidence* thesis. There are three main general worries about this thesis. The first is *metaphysical*. The basic worry is that ethical properties (and our beliefs about these properties) are *weird*, so that ethical propositions are just not the right sort of thing to count as evidence for and against empirical conclusions. The second worry is *epistemological*. The point is that, even if ethical propositions could legitimately count as empirical evidence, they could not be *relied upon* as evidence for or against empirical conclusions, because our epistemic justification for believing ethical propositions is not sufficiently independent of our justification for believing these empirical conclusions. The final worry is just that there are not many clear *examples* of cases where we can legitimately rely upon ethical propositions as evidence for empirical conclusions (but there are plenty of cases where it appears to be irrational to do so).

§1 counters the metaphysical worry by showing that, on extremely minimal (and widely shared) assumptions about the metaphysics of ethical properties and the nature of ethical belief, ethical propositions can often meet the standard criteria for counting as empirical evidence according to our best general theories of evidential support.

§2 responds to the epistemological worry by showing that, if we can gain epistemically justified ethical beliefs by means of testimony, then ethical beliefs can be justified sufficiently independently of beliefs about empirical conclusions in order to be rationally relied upon as evidence for and against these conclusions.

exceptions to this trend, especially from feminist and pragmatist philosophers of science who have argued that value-judgements *can* play an evidential or quasi-evidential role in scientific inquiry (e.g., White 1981, Nelson 1990, Campbell 1998, Anderson 2004, Clough 2003, 2011, Brown 2020). The aim of this paper is to build on their work by giving a clearer and more systematic treatment of the metaphysics and epistemology underlying this position.

Outside of the values-in-science debate, the question of whether ethical propositions can be used as evidence for empirical conclusions has also recently cropped up in metaethics, with Barber (2013) and Basl and Coons (2017) each trying to use a negative answer to this question in order to motivate a controversial result about ethical metaphysics or epistemology (see §2 below). However, it is worth noting that in other areas of this literature, it is assumed to be true without much argument that ethical propositions can legitimately be used as evidence for empirical conclusions. For instance, this features as a key premise in Dorr's (2002) influential argument against certain versions of metaethical non-cognitivism, and in Podgorski's (2020) recent argument about the ethical significance of moral uncertainty.

§3 rounds off the discussion by outlining several examples where it *is* plausible that we can legitimately use an ethical proposition which we justifiably believe as evidence for or against an empirical conclusion.

The goal of this paper is not to show that one *must* accept the Ethical Evidence thesis. Rather, the aim is to show that (i) the legitimacy of relying on ethical propositions as empirical evidence follows from uncontroversial assumptions about ethical metaphysics and epistemology, and (ii) there are cases where it is quite plausible that we can legitimately rely upon ethical propositions which we justifiably believe as empirical evidence, therefore, (iii) there is no good reason to reject these uncontroversial assumptions. In other words, we should see the legitimacy of evidentially relying upon ethical propositions as the *default position* in this debate.

1. Metaphysics

This section discusses the metaphysical objection to the Ethical Evidence thesis, which trades off the idea that ethical properties are metaphysically *different in kind*, in some important respect, from empirical properties. Additionally, beliefs about ethical propositions appear to be different in kind from beliefs about empirical propositions, insofar as ethical beliefs appear to have *conative* as well as cognitive content.⁴ Therefore, ethical propositions might not be the *right sort of thing* to count as evidence for and against empirical conclusions.

In the longstanding philosophical debate about the role of value-judgements in scientific inquiry, something like this general metaphysical worry has historically led many philosophers and scientists to reject the Ethical Evidence thesis. Although this argument is almost never spelled out in detail, the idea seems to be something like the following: (i) value-judgements (i.e., ethical beliefs) are a desire-like attitude; (ii) allowing one's desires to influence whether one believes a particular empirical conclusion would

⁴ To clarify: what I call the 'metaphysical' objection to Ethical Evidence involves both a worry about the metaphysics of ethical properties and about the content of ethical beliefs. This latter worry is not strictly about ethical metaphysics, but the metaphysics of ethical properties and the content of ethical beliefs are typically discussed together in the metaethical literature, so I lump them together here.

amount to a kind of *wishful thinking*; (iii) engaging in wishful thinking is always epistemically impermissible; therefore (iv) one can never rely upon their value-judgements (or the ethical propositions which are the content of these judgements) as evidence for or against empirical conclusions.

However, although this kind of wishful-thinking objection to Ethical Evidence has been very popular in the debate about the role of values in scientific inquiry (e.g., Geertz 1990:19, Haack 1993:42, Gross and Levitt 1994: 162, Douglas 2009:87, Brown 2013, Elliott 2017:11), it does not hold up to scrutiny. Even if ethical beliefs have conative content or are desire-like in some respect,⁵ this does not mean that relying upon ethical beliefs as evidence for empirical conclusions would involve the same kind of (epistemically impermissible) transition which is involved in stereotypical cases of wishful thinking. In particular, those who endorse the wishful thinking objection seem to have the following inference in mind as the archetypal example of using ethical propositions to influence empirical beliefs:

Wishful Thinking:

It would be ethically/morally/politically good if P were the case.

Therefore, P is the case.

Although this inference does involve a kind of wishful thinking, there is no reason to think that this, in itself, raises any serious problems for the general Ethical Evidence thesis. That is, this thesis says only that ethical propositions can be relied upon as evidence for and against empirical conclusion as long as they meet the standard criteria for counting as evidence for or against these empirical conclusions according to our best general theories of evidential support. However, since the mere fact *that it would be good if P were the case* is not generally good evidence for the conclusion *that P actually is the case* according to any plausible theory of evidential support, then the Wishful Thinking inference cannot be representative of the kinds of inference which would be licensed by the truth of the Ethical Evidence thesis.

⁵ See Bedke (2017) for an introduction and overview of the debate about the conative content of ethical beliefs.

So, the general metaphysical objection to the Ethical Evidence thesis is not best understood in terms of a popular worry about wishful thinking. However, the above discussion does motivate a better way of understanding the metaphysical objection. The Ethical Evidence thesis says that ethical propositions can be legitimately relied upon as evidence as long as they meet the standard criteria according to our best theories of evidential support. However, the weirdness of ethical metaphysics might make it the case that *it is impossible for ethical propositions to meet these criteria*.⁶ For instance, if ethical propositions are not truth-conditional (e.g., according to metaethical emotivism), or if the truth and falsity of ethical propositions does not co-vary at all with the truth and falsity of propositions about the empirical world (e.g., according to an extreme form of metaethical relativism or non-naturalism), then these propositions would simply not be capable of standing in the evidential support relation to any empirical conclusion. In other words, if one endorsed an *extremely* anti-realist metaethical theory (or perhaps an extreme form of non-naturalism), then one might be able to argue that ethical propositions cannot meet the standard criteria for counting as evidence for and against empirical conclusions, and so they cannot be relied upon as evidence for believing or rejecting these empirical conclusions.

However, this kind of metaethical position is way out of line with the standard and widely-shared assumptions of the current metaethical debate. In particular, the following four assumptions, which are accepted by practically all metaethicists working today, are sufficient to entail that ethical propositions are the right sort of thing to count as evidence for and against particular empirical conclusions:

1. *Minimal Cognitivism*: We have beliefs (or belief-like states) with ethical propositions as content, and these propositions and beliefs are typically either true or false (at least in some minimal sense).⁷
(E.g., it is true that Spencer cut the cake unfairly).

⁶ Technically, this objection would render the Ethical Evidence thesis trivially true rather than false.

⁷ It is worth emphasising that even metaethical 'non-cognitivists' almost universally accept this kind of minimal cognitivism and truth-aptness. See, e.g., Bedke (2017:297-8).

2. *Ethical Supervenience*: The truth or falsity of any given (contingent) ethical proposition systematically modally covaries with the truth or falsity of some particular set of empirical propositions.⁸
(E.g., the truth of the proposition that Spencer cut the cake unfairly systematically covaries with the truth of the proposition that Spencer cut the cake unevenly.)
3. *Justified Beliefs about Covariance Patterns*: We have *some* justified beliefs about the patterns of covariance between the ethical and the empirical.
(E.g., I am justified in believing that, all other things equal, if someone cutting a cake for two people cuts it unevenly, then they cut it unfairly.)
4. *Explanatory Connection*: The covariance patterns between the ethical and empirical are mediated by a (perhaps indirect) explanatory connection.
(E.g., Metaethical realists say that the fact that Spencer cut the cake unfairly is (partly) grounded in the empirical fact that he cut the cake unevenly. Anti-realists say something more complicated like, this ethical fact is grounded in facts about our evaluative attitudes and dispositions, and these facts about our minds are themselves systematically explanatorily connected to empirical facts about the unevenness of the cake).⁹

How do these assumptions make it the case that ethical propositions meet the standard criteria for counting as evidence in favour of empirical conclusions? First, let's take the standard *probabilistic account* of evidential support defended by Timothy Williamson (2000), among many others. Very roughly, on this account, a proposition P is evidence for a particular agent to believe another proposition Q just in case the probability of Q being true conditional on P being true is greater than the unconditional probability of Q being

⁸ McPherson (2012:205) notes that 'the supervenience of the ethical has been a rare locus of near-consensus in metaethics'.

⁹ For explanatory connections in realist theories, see Rosen (2017). Anti-realist philosophers are usually less explicit about the (indirect) explanatory connection between the empirical and ethical facts in abstract discussion, but it is easy to extract a commitment to this assumption from their specific theories. For example, on Gibbard's (2003) influential expressivist metaethical theory, the explanation for why a particular contingent ethical statement is true is (roughly) that (i) a relevant bunch of empirical facts P are true *and* (ii) the agent (relative to whom the ethical fact can be treated as a true) has the relevant set of 'planning attitudes'. Even though the explanatory connection is indirect, the truth of a particular empirical proposition P can nevertheless feature as *part* of the explanation for why a particular ethical proposition E is true.

true. In other words, for an agent S, two propositions P and Q and a subjective probability function $\Pr()$:

P is evidence for S to believe Q *if and only if* $\Pr(Q|P) > \Pr(Q)$

(adapted from Williamson 2000:187)

We can understand $\Pr(Q|P)$ as being determined either by the objective probability of Q's being true conditional on P's being true (i.e., by the objective covariance patterns between P's truth and Q's truth) or as being determined by the subjective probability of Q's being true conditional on P's being true (i.e., by S's justified beliefs about these covariance patterns). It doesn't matter for the purposes of the present discussion. The point is just that on either of these options, the assumptions above ensure that ethical propositions are metaphysically the right sort of thing to count as evidence in favour of empirical conclusions. For instance, I may start out with a low confidence that a particular empirical conclusion is true; e.g., I initially think it's unlikely that Spencer will cut the cake unevenly. But regardless of whether the relevant probabilities are being determined objectively or relative to my justified beliefs, the conditional probability of Spencer having cut the cake unevenly conditional on it being true that Spencer cut the cake unfairly is pretty high; so if this latter ethical proposition became part of my total body of evidence, this evidence would tell me to increase my confidence that Spencer cut the cake unevenly. And, more generally, given that the truth of ethical propositions generally systematically covaries with the truth of empirical propositions (assumption 2) and we plausibly often have justified beliefs about this covariance (assumption 3), on the probabilistic account of evidential support, it looks like ethical propositions will *often* be the right sort of thing to count as evidence for and against empirical conclusions.

The other popular general account of evidential support is *explanationism*, according to which one proposition is evidence for another proposition just in case there is (or appears to be) the right kind of explanatory connection between these propositions (e.g., McCain 2013). In particular, explanationists generally argue that the following is a sufficient condition for evidential support:

P is evidence for S to believe Q if Q's being true is part of the Best Explanation (on S's total background evidence) for why P is true.

(adapted from McCain 2013:302)¹⁰

The Best Explanation is the explanation which scores highest (and perhaps also high enough) according to the explanatory virtues of simplicity, coherence, depth, etc (for an overview, see Lipton 2003). Although there has historically been disagreement over whether ethical propositions can ever feature in the Best Explanation of empirical truths (e.g., Majors 2007), the *contrary* is uncontroversial. It should be clear (given assumption 4 above) that an empirical proposition can feature as part of the Best Explanation (on one's total background evidence) for why an ethical proposition is true; for example, the fact that Spencer cut the cake *unevenly* can feature as part of the Best Explanation for why it is true that Spencer cut the cake *unfairly*. Therefore, ethical propositions (as the propositions whose truth needs to be explained) are often the right sort of thing, metaphysically, to be evidence in favour of empirical conclusions (as the propositions whose truth does the explaining).

To reiterate: as long as we accept these four assumptions about ethical metaphysics (which are compatible with a wide array of metaethical theories, including reasonably extreme forms of non-naturalism and anti-realism), then ethical propositions can meet the standard criteria for counting as evidence for or against empirical conclusions.¹¹ In other words, even though the metaphysics of ethical properties (or the content of ethical beliefs) may be quite different from the metaphysics of empirical properties (or the content of empirical beliefs), these four assumptions entail that ethical propositions are just *not weird enough* to be the wrong kind of things to count as evidence for empirical

¹⁰ On McCain's account of evidential support, evidence is constituted by *appearances* rather than propositions. As I discussed in fn.2, this does not substantively change the discussion in this paper. Additionally, McCain thinks that the explanandum of this Best Explanation is not just *why P is true*, but rather *why P is possessed as evidence* (i.e., why it seems to one as if P, or why one has a justified belief that P, etc.). This does not matter for the current discussion: if an empirical proposition can feature in the Best Explanation of why a particular ethical proposition is true, then it is uncontroversial that this empirical proposition can also feature as part of the Best Explanation of why it seems to someone as if that ethical proposition is true.

¹¹ Basl and Coons (2017:162-7) argue for a similar conclusion from the opposite direction. That is, while I argue from first principles (i.e., assumptions which are widely shared in the current metaethical debate), they argue via elimination from existing categories of metaethical theories. See also Brown (2020:89-100, 163-165), who argues that on any 'modest' cognitivist metaethical theory, our beliefs about ethical propositions (or value-judgments more broadly) will be the right sort of thing to play an evidential role in scientific inquiry.

conclusions. And although it is beyond the scope of this paper to give any specific defence of these assumptions, it should be clear that rejecting them would come at a heavy price. The big-picture point at the present stage of the discussion is just that there is no reason to think that this price is worth paying in order to reject the Ethical Evidence thesis.

2. Epistemology

However, even if ethical propositions can meet the standard criteria for counting as evidence for or against specific empirical conclusions, there may be other reasons why ethical propositions cannot be *relied upon* as one's evidential basis for believing these conclusions. In particular, Alex Barber (2013) and John Basl and Christian Coons (2017) have recently argued that there are epistemological problems with relying upon ethical propositions as evidence for or against empirical conclusions.

Barber, Basl and Coons do not set out to *argue* that the Ethical Evidence thesis is false. Rather, they begin by assuming that ethical propositions cannot be relied upon as evidence, in order to argue for more controversial conclusions about ethical epistemology.¹² They motivate this assumption (i.e., that Ethical Evidence is false) by pointing to specific examples where it appears to be clear that inferring from an ethical or normative premise to an empirical conclusion would be epistemically irrational, such as the following:

Bed Bomb:

¹² It is worth noting that Basl and Coons (and to some degree Barber) are only concerned with the Ethical Evidence thesis insofar as it bears on the question, 'why doesn't *moral reasoning* serve as a tool for learning about the physical world?' (Basl and Coons 2017:161, emphasis added). In other words, these authors are more concerned with why it is impermissible to go about researching or *investigating* empirical reality by means of investigating ethical truths. However, (i) their arguments against "moral science" still bear directly on the Ethical Evidence thesis, so I focus on this aspect of their discussion in the present section; (ii) on a wider, more social interpretation of what 'investigation' amounts to, it is plausible that *seeking out reliable testimony* can be a means of investigation – and if so, then this section of the paper should also be interpreted as arguing that we *can* investigate the empirical world by means of ethical inquiry and investigation; and finally, (iii) although it is beyond the scope of this paper to defend this stronger thesis, Brown (2020) argues that we can (and should) do scientific investigation and inquiry which is epistemically influenced by our moral reasoning and imagination – which if correct might also entail that a form of "moral science" is epistemically permissible. Thanks to a reviewer for pressing me to say more about this.

I ought to get out of bed.

I oughtn't get out of bed if a weight-sensitive bomb is under my bed.

Therefore:

There is no weight-sensitive bomb under my bed.

(Adapted from Basl and Coons 2017:174).¹³

This inference is deductively valid, so it seems like the ethical/normative premise is the right sort of thing to count as evidence for the empirical conclusion. So why can't a (typical) agent gain epistemic justification for believing the empirical conclusion by means of inferring from the ethical premise? Barber (2013) and Basl and Coons (2017) argue that this comes as a result of the distinctive epistemology of ethical propositions.¹⁴

The basic strategy here is to appeal to the extensive epistemological literature on the *failure of transmission of justification* (e.g., Wright 2002, Tucker 2010, Moretti and Piazza 2013).¹⁵ This literature is concerned with the general class of deductive inferences where the premises entail that the conclusion is true but making the inference does not grant an agent any epistemic justification for believing the conclusion. For instance, the following is often cited as an example of an inference where there is transmission-failure:

Zebra:

¹³ Barber (2013:634) and Basl and Coons (2017:161, 168) give several other examples of ethical-empirical inferences which also appear to be epistemically irrational. The discussion below does not specifically address all of these examples, and one might think that some of them fail to transmit justification for reasons which are more complicated than those I discuss here. This is probably true, but the basic point remains that *even if* ethical propositions can generally be relied upon as empirical evidence, there would *still* be a wide range of cases where inferring from an ethical proposition to an empirical conclusion fails to give an agent epistemic justification for believing the conclusion, but the explanation for this has to do with (perhaps quite complicated) *general* epistemological principles governing the transmission of justification rather than anything to do with the ethical nature of the premises.

¹⁴ Barber (2013:637-9) also entertains the possibility that this inference fails to transmit justification because some extreme form of metaethical anti-realism is true. However, Basl and Coons (2017:162-7) convincingly argue that Barber's discussion here is largely misguided.

¹⁵ Neither Barber nor Basl and Coons explicitly reference this extensive literature in their discussion, although it is clear that this is precisely the epistemological problem with which they are concerned. In what follows, I aim to clearly reconstruct the best version of their arguments, rather than sticking to strict exegesis.

This animal is a zebra.

If this animal is a zebra, then it is not a cleverly-disguised mule.

Therefore:

This animal is not a cleverly-disguised mule.

Most epistemologists think that making this inference would not (typically) give an agent any new or additional justification for believing that the animal is not a cleverly-disguised mule.¹⁶ There is some disagreement over the exact diagnosis of *why* justification is not transmitted in this kind of inference, but the basic thought is that one's epistemic justification for believing *that this animal is a zebra* (usually) wholly depends on having prior epistemic justification for believing *that this animal looks like a zebra*. However, this proposition is no better evidence for the conclusion *that this animal is a zebra* than for the conclusion *that this animal is a mule cleverly disguised to look like a zebra*,¹⁷ so even though these two conclusions are inconsistent, one is not in any position to make an inference from the truth of one of them (i.e., *that this animal is a zebra*) and thereby gain justification for believing the falsity of the other (i.e., *that this animal is not a cleverly-disguised mule*).

Epistemologists have sought to extract a more general principle from this kind of diagnosis of the case, which is standardly given along the following lines:

Failure: If one's antecedent epistemic justification for believing P is not (at least partly) *counterfactually independent* of one's antecedent justification for believing Q, then one cannot infer from the premise

¹⁶ A prominent exception is Tucker (2010), who argues that one actually *can* gain (doxastic) justification for believing that the animal is not a cleverly-disguised mule by believing this on the basis of one's belief that the animal is a zebra. More generally, it is worth stressing that even those epistemologists who think that an agent cannot *usually* gain epistemic justification for believing the conclusion by making this inference *should* think that there are certain special cases where justification *does* transmit. In particular, if one received *reliable testimony* from a zookeeper that this animal is a zebra, then it seems that one could infer from the fact that it is a zebra and gain epistemic justification in the conclusion that it is not a cleverly-disguised mule.

¹⁷ At least, the proposition *that this animal looks like a zebra* is no better evidence for one conclusion over the other *once the sceptical possibility is raised* (i.e., once it is relevant to consider the possibility that the animal is a cleverly-disguised mule). See J. Brown (2016).

P to the conclusion Q and thereby gain epistemic justification for believing Q. (Moretti and Piazza 2013:2488)

For what it's worth, this principle is probably too strong.¹⁸ However, it appears to be exactly the kind of claim that Barber (2013:644-5) and Basl and Coons (2017:177-8) rely upon in their argument, so I will assume that it is true in what follows. In any case, if the discussion below succeeds in showing that inferences from ethical premises to empirical conclusions *can* transmit justification while assuming an unnecessarily strong principle for transmission-failure, then the argument will also go through on a weaker principle; so there is no harm in assuming that *Failure* is correct for the purpose of the present argument.

Returning to the specific example above, the point is that making the inference from the ethical premise *that I ought to get out of bed* to the conclusion *that there is no weight-sensitive bomb under my bed* fails to transmit justification for the reason described by *Failure*. That is, one's antecedent justification for believing that they ought to get out of bed (normally) wholly counterfactually depends on their antecedent justification for believing that there is no weight-sensitive bomb under their bed. Therefore, according to the *Failure* principle, inferring from this ethical premise cannot give an agent any new or additional epistemic justification for believing the conclusion.

This does appear to give a broadly correct diagnosis of what goes wrong in the specific inference above. But this is not obviously a big deal: if ethical propositions can generally be relied upon as evidence for and against empirical conclusions, then it would not be surprising that there are *some* possible inferences from ethical premises to empirical conclusions which fail to transmit justification, for general epistemological reasons rather than reasons related to the fact that the premise happens to be an ethical proposition.

However, Barber and Basl and Coons want to say something much stronger here. They want to generalise from this specific example to the general class of *all* (putatively sound)

¹⁸ Moretti and Piazza (2013) give a sophisticated counterexample to this principle involving probabilistic outcomes. However, a more obvious problem is that this principle appears to entail the highly counterintuitive result that, whenever one infers from P to Q and then from Q to R, this second inference *always* fails to transmit justification because one's (propositional) epistemic justification for believing both Q and R wholly depend on one's antecedent justification for believing P. See Tucker (2010:512-4) for further discussion.

deductive and inductive inferences which feature an ethical proposition as a premise and an empirical proposition as the conclusion.¹⁹ The idea is to say that one's epistemic justification for believing *any* ethical proposition wholly counterfactually depends on their antecedent justification for believing *all* of the empirical conclusions in favour of which this ethical proposition would be good evidence; so that if any agent ever tried to infer from an ethical proposition to an empirical conclusion for which the former is apparently good evidence, one could never gain justification for believing the latter. That was a bit of a mouthful, but to put it simply: ethical propositions can (nominally, at least) count as good evidence for empirical propositions, but because justified beliefs about ethical propositions are always epistemically mediated by justified beliefs about empirical propositions, one can never make a good inference from an ethical premise and thereby gain justification for believing an empirical conclusion (Barber 2013:645-9, Basl and Coons 2017:179-80).

It is certainly plausible that there is a general class of cases where one can only gain justification for believing an ethical proposition by means of relying on their prior justification for empirical beliefs about the facts which would ground the truth of this ethical proposition. For instance, Basl and Coons (2017:179) argue that cases of direct ethical knowledge gained via exercising one's capacity for *moral perception* plausibly require that one's access to contingent ethical facts be mediated by perceptual access to the empirical facts which ground these facts.²⁰ For instance, in the case described in my anecdote at the beginning of this paper, when Emily forms a justified ethical belief about whether Spencer cut the cake fairly, her justification for believing that the cake is cut unfairly wholly counterfactually depends on her antecedently-justified perceptual belief that the cake was cut unevenly (combined with a bridge-assumption linking unevenness

¹⁹ Or, at least, for most cases. For instance, Basl and Coons (2017:180-4) argue that there is a different reason why inferences from ethical premises to empirical conclusions fail to transmit justification when these ethical and empirical propositions purport to describe necessary truths or law-like statements. We can put this class of inferences aside for the purpose of the present discussion.

²⁰ I follow Basl and Coons in using the phrase, 'moral perception', slightly out of line with the general philosophical literature here. Strictly speaking, what I have said above amounts to the rejection of the thesis that we gain ethical knowledge through a distinctively *moral perception*, and instead, this kind of direct ethical knowledge comes via making an inference from one's justified empirical beliefs. This position is not uncontroversial; see Werner (2020) for a recent overview of the debate around moral perception.

and unfairness in this context). Therefore, although the fact that the cake was cut unfairly is objective evidence that the cake was cut unevenly, Emily *herself* could not use this evidence to increase her justification for believing that the cake was cut unevenly.

However, it is plausible that there is another general class of cases where one's ethical beliefs *can* be justified broadly independently of one's beliefs about the empirical propositions for which these ethical propositions would be evidence. That is, it is plausible that many of our ethical beliefs are justified by means of *reliable testimony*. This can either be direct testimony from particular individuals (as when Emily told me *that Spencer cut the cake unfairly*) or a more diffuse kind of social testimony (as when we come to know *that actual patterns of racial inequality are unjust*). In these cases, the justification of one's ethical belief does *not* (wholly) counterfactually depend on their antecedent justification for believing, for instance, the empirical proposition(s) whose truth would ground the truth of the relevant ethical proposition. Therefore, the *Failure* principle does not rule out that one *could* infer from a testimonially-justified belief in an ethical proposition to the conclusion that a particular empirical proposition is true, and thereby gain justification for believing this empirical conclusion.²¹ In other words, when our ethical beliefs are justified (at least partly) via reliable testimony, justification *can* be transmitted from ethical premises to empirical conclusions, and this is plausibly exactly

²¹ At this stage, it is instructive to consider the following objection. If one accepts a strong form of reductionism about testimonial justification, then one thinks (roughly) that one can gain epistemic justification for a testimonial belief that X only if they make a quasi-inference from the fact *that such-and-such said X* (Lackey 2006:160-3). So, it seems that if one gained a belief in an ethical proposition via testimony, one's epistemic justification for believing this proposition would wholly depend on one's prior epistemic justification for believing an empirical proposition about the testifier's assertion. So, if one embraces this controversial reductionist view about testimonial justification, then one's justification for believing ethical propositions will still depend on one's justification for believing *some* empirical proposition. Is this a problem for my argument? No. First, it is not clear that, just because one's justification for believing a particular ethical proposition E depends on one's prior justification for believing the particular empirical proposition *that S testified that E*, one is therefore not in a position to use E as evidence for a *different* empirical conclusion (i.e., one for which the fact *that S testified that E* would not be good evidence anyway).

Second, and more importantly, *even if* this objection showed that ethical propositions learned via testimony are not sufficiently independently justified in order to be relied upon as evidence, the argument would also extend much further. That is, this objection would show that one can never use *any* proposition P which one learns via testimony as evidence, since its evidential relevance is always 'screened off' by the dependence on one's prior justification for believing *that so-and-so asserted that P*. Even if one were to embrace this implausible conclusion, the *comparative* conclusion would still remain that *there is no relevant* difference between the evidential role played by ethical propositions and non-ethical propositions which are learned via testimony (c.f., n.2 above).

what happened when I made the (successful) inference from the premise *that Spencer cut the cake unfairly* to the conclusion *that Spencer cut the cake unevenly*.

Barber (2013:644-5) and Basl and Coons (2017:178-9) do recognise the possibility that, if one *could* get reliable testimony about contingent ethical facts, then one could use these ethical facts as evidence for and against empirical conclusions. However, they don't make a big deal about this. On the one hand, Barber appears to suggest that reliable testimony about ethical facts can only come from supernatural beings like an infallible oracle or God (Barber 2013:645). This point is exaggerated. It is highly plausible that there are many cases of reliable ethical testimony from fallible sources; for instance, Emily's testimony *that Spencer cut the cake unfairly* is a clear example. Nevertheless, some philosophers *have* thought that there are distinctive problems about relying on ethical testimony (e.g., Hills 2009, Crisp 2014, Mogensen 2017; though for optimism about ethical testimony, see, e.g., Sliwa 2012, Groll and Decker 2014, Wiland 2021). For instance, there appears to be something odd with simply accepting someone's testimony *that eating meat is morally wrong*, and therefore becoming a vegetarian. If these *ethical testimony pessimists* are right, then it might appear that reliable testimony is not a valid source of epistemic justification for believing ethical propositions.

However, even if there is something problematic about accepting ethical testimony, ethical testimony pessimists generally argue that these are *practical* problems rather than epistemological ones. That is, there is broad consensus that *even if* there are problems with relying upon ethical testimony in practical reasoning or action (e.g., due to a lack of autonomy or acting without moral understanding), agents can still readily gain epistemically justified ethical beliefs – and even ethical knowledge – via testimony (Hills 2013:553-4, McShane 2021, Wiland 2021:28-47). Therefore, even if it would be wrong or improper to *act* upon ethical beliefs which one gained via reliable testimony, there is no reason to think that these beliefs would lack the epistemic justification necessary for possessing the contents as evidence.²²

²² One way to see this is to consider the case of a testimonially justified ethical belief being used as evidence for *other* ethical propositions. For instance, suppose someone tells me that it's wrong to eat meat, and I accept this testimony, thereby forming the relevant ethical belief in this proposition. Since I already (independently) know some relevant empirical propositions (e.g., the meat-production industry causes animals a lot of pain), I can plausibly use the ethical proposition which I have just learned to (justifiably) infer the ethical conclusion *that this pain morally outweighs the pleasure that humans get from eating meat*. In other words, when I learn

On the other hand, Basl and Coons (2017:178-9) suggest that it matters in cases of ethical testimony that, even though the *recipient* of reliable testimony need not have antecedent justification for believing the relevant empirical propositions, the *testifier* (or the original testifier in a long chain of testifiers) must have had antecedent justification for believing these empirical propositions. The idea appears to be that, as long as there is epistemic dependence on antecedently justified empirical beliefs *somewhere* in the testimonial chain, then no recipient of this testimony can rationally rely on the ethical proposition as evidence that these empirical propositions are true.

However, it is unclear why this kind of indirect dependence on *someone else's* justification for believing empirical conclusions should prevent one from rationally relying upon the ethical propositions which they justifiably believe as evidence for these empirical conclusions. For example, this clearly did not make a difference to my epistemic situation when I received reliable testimony from Emily *that Spencer cut the cake unfairly*; even though *she* could not use this ethical proposition as evidence to gain justification for believing that Spencer cut the cake unevenly, upon receiving her reliable testimony, *I myself* was clearly in a position to use this ethical proposition as evidence.

Perhaps the objection here is more subtle, and is concerned not so much with *individuals* using ethical propositions as evidence for particular empirical conclusions, but rather with the use of ethical propositions as evidence *collectively*. For instance, if ethical propositions could legitimately count as evidence for empirical conclusions, then a radical implication might be that the *scientific community* should change its practice such that our best scientific theories should be weighed against our justified ethical beliefs. But this would require *more* than the simple condition that some individual's ethical belief be justified independently of their relevant empirical beliefs; in particular, it would require that the relevant ethical belief of the scientific community (*taken as a whole*) be justified independently of the relevant empirical beliefs of the community (taken as a whole). However, if the ethical beliefs of *some* of the testifiers in the relevant testimonial chain depend for their justification on prior justification for believing the relevant empirical

an ethical proposition via testimony, I am plausibly in a position to use this ethical proposition as evidence for other ethical conclusions, in order to gain ethical understanding (c.f., Wiland 2021:56-7). Likewise, when I learn an ethical proposition via testimony and I know the relevant ethical bridge-principles but not the relevant empirical facts, I can plausibly legitimately use this ethical proposition as evidence in favour of the relevant empirical conclusions.

conclusions, then it may well be the case that the relevant ethical belief of the scientific community depends for its justification on antecedently-justified collective empirical beliefs. Therefore, the scientific community could not collectively rely on this ethical proposition as evidence for that empirical conclusion, even if some individuals in this community could individually do so.

This is an interesting argument, and although it does not show that ethical propositions cannot in general be relied upon as evidence for empirical conclusions, it does take the sting out of this conclusion. Effectively, this line of argument seeks to show that, even if it is true that ethical propositions can count as empirical evidence, this conclusion may not make any difference to our practice of scientific inquiry. Although this objection does not pose a direct challenge to the Ethical Evidence thesis, it is worth briefly mentioning two plausible strategies for responding to it.

First, it is worth emphasising that many individuals have been, and continue to be, systematically excluded from the scientific community. Moreover, it is *specifically* these marginalised individuals who are likely to occupy the right kind of epistemic standpoint to have justified beliefs about many ethical propositions;²³ for instance, women have historically been both excluded from the (social) scientific community and have also been in distinctive position to gain the justified ethical belief that, e.g., *divorce is not inherently morally bad* (c.f., Anderson 2004:12). Therefore, if any of these marginalised individuals passed on their justified ethical belief to members of the scientific community through reliable testimony, then (because of this person's exclusion from this epistemic group) the scientific community's justification for (collectively) believing this ethical proposition would not wholly depend on their antecedent justification for (collectively) believing the empirical propositions for which this ethical proposition would be evidence.

Second, in the background of Basl and Coons's argument is a picture of ethical testimony as a *chain of individuals*, the first of whom forms an ethical belief (e.g., via 'moral perception') and whose epistemic justification for holding this ethical belief wholly

²³ For a recent overview of the philosophical literature on standpoint epistemology, see Toole (2021). Standpoint theorists tend not to explicitly state that some of the truths which can be known from an marginalised standpoint have ethical content. However, this follows readily from some other claims of standpoint epistemology, in particular, that oppressed people are in a better position to understand the features of their oppression (which plausibly includes the normative features of this oppression).

counterfactually depends on their antecedent justification for holding a number of empirical beliefs; then this justified ethical belief is passed from testifier to testifier right down the linear chain. However, although this simple picture is useful for some kinds of analysis of how testimonial justification works, the reality is much more complicated. Just like the social epistemology of scientific knowledge, the social epistemology of ethics is often extremely diffuse and decentralised (Anderson 2016). Take, for instance, my (justified) belief *that actual persistent patterns of income and wealth inequality in highly developed countries are unjust* (for instance, that the contingent, actual state of affairs where Black communities in the USA are generally poorer and have less social mobility than White communities is unjust). This belief was justified through diffuse testimony from my ethical community; that is, I do not have (or at least, do not need to have) independently justified beliefs about the empirical propositions whose truth would potentially ground the truth of this ethical proposition. Instead, this belief comes from a number of testimonial sources, as a result of being embedded in a particular social position and being influenced by a number of individuals who share this belief (but whose beliefs themselves were also justified through a similar diffuse testimonial process). At no level of this process is there an individual who has independently-justified beliefs about *all* of the relevant empirical propositions, and infers that actual inequality is unjust. Instead, there is diffuse, feedback-looping testimonial input all the way down.

It is beyond the scope of this paper to say anything distinctive or insightful about how this diffuse process of ethical social epistemology works, or what the original sources of epistemic justification in this process are. The big-picture point for now is just that, when the collective production of ethical knowledge in our communities systematically tracks the (contingent) ethical features of reality through a diffuse and decentralised social epistemological process, it is not plausible that the justification of this community's ethical beliefs wholly counterfactually depends on their collective antecedent justification for believing the empirical propositions whose truth would ground the truth of the relevant ethical propositions.²⁴

²⁴ It is worth adding to this that, as a sociological fact, the production of ethical knowledge *often* operates largely independently of the production of empirical knowledge by the scientific community. For instance, Anderson (2016) documents how the ethical knowledge *that the actual, historical practice of slavery was morally impermissible* was generated primarily by people under risk of being enslaved themselves, while the European scientific and intellectual community continued to believe both that slavery was permissible and that Black slaves were

To sum up: if it were impossible to gain epistemic justification for believing ethical propositions which is (at least partly) counterfactually independent of our antecedent justification for believing the empirical propositions for which these ethical propositions count as evidence, then we could not rationally *rely upon* these ethical propositions as our evidential basis for believing these empirical conclusions. However, it is plausible that we can gain epistemic justification for believing at least some ethical propositions by means of reliable testimony, and if so, then these ethical beliefs can be sufficiently independently justified to provide an appropriate evidential basis for believing particular empirical conclusions. Some philosophers have raised worries about the possibility of gaining epistemic justification by means of ethical testimony, either for individuals or for epistemic groups or communities. However, there is good reason to think that (i) any problems with ethical testimony for individuals are practical rather than epistemic in nature, and (ii) there are several ways in which groups can gain epistemic justification for ethical beliefs by means of reliable testimony (either because some testifiers are excluded from these groups or because of the diffuse nature of the testimony in question). In sum, the *default position* should be that we sometimes have the right kind of epistemic justification for believing ethical propositions in order to rationally rely upon these propositions as evidence for particular empirical conclusions.

3. Examples

At this stage of the paper, I take myself to have shown that the Ethical Evidence thesis follows from standard assumptions about ethical metaphysics and epistemology. This result can be seen as a strong consideration in favour of the Ethical Evidence thesis, but it can also be seen as a *reductio ad absurdum* of at least one of these standard assumptions. That is, if we had independent reason to think that it is always epistemically impermissible to rely on an ethical proposition as evidence in favour of an empirical conclusion – if, for instance, all of the examples of ethical-empirical inferences which we

intellectually and emotionally inferior to free Whites. When attitudes to the moral permissibility of slavery changed widely in the 19th century, this change was prompted by engagement with the lived experience of the oppressed and not by revised scientific beliefs about racial intellectual equality amongst European scientists.

can think of involve an intuitively epistemically impermissible transition – then we should reject one or more of the assumptions which collectively entail Ethical Evidence.

As I briefly mentioned at the beginning of §2, something like this *reductio ad absurdum* argument is the strategy pursued by Barber (2013) and Basl and Coons (2017), who argue on the basis of the apparent epistemic impermissibility of particular deductive inferences from ethical premises to empirical conclusions (like *Bed Bomb*, above), for the conclusion that we need to either accept an (implausibly) extreme anti-realist metaethical theory (Barber 2013:637-9) or embrace an extreme picture of ethical epistemology according to which one cannot gain epistemic justification for believing an ethical proposition without having prior justification for believing *all* of the empirical propositions for which this ethical proposition would count as evidence (Barber 2013:645-51; Basl and Coons 2017:174-80).

I argued in the last section that the epistemic impermissibility of inferences like *Bed Bomb* can be explained using general epistemological resources, and in particular, the resources from the large epistemological literature on the failure of justification-transmission. So, there is no reason to think that the epistemic impermissibility of making these inferences comes as a result of the mere fact that an ethical proposition is being used as evidence for an empirical conclusion. However, this is not sufficient to refute the *reductio ad absurdum* argument unless we can also find plausible examples where it *is* epistemically permissible (and even epistemically virtuous) to use an ethical proposition as evidence for or against an empirical conclusion, in exactly the way that the Ethical Evidence thesis would predict. This section of the paper aims to outline several such examples, both for the case of an individual believer and also for the case of collaborative social scientific practice.

It is actually quite easy to find epistemically permissible ethical-empirical inferences for individual believers. I outlined one such example at the very beginning of this paper, when I described how I inferred from the ethical proposition *that a piece of cake was cut unfairly* (which I learned via reliable testimony), and came to have a justified true belief in the empirical conclusion *that the cake was cut unevenly*. This example can be abstracted into the following general schema:

Unfair to Unequal

X is unfair (in some respect).

Therefore: X is (probably) unequal (in some respect).

This inference is not deductive (although it could be made into a deductive inference by adding as a second premise the assumption that *if X is unfair in some respect, then X is unequal in some respect*). However, a deductive inference is not necessary here; it is enough that the premise is good but fallible evidence for the conclusion. That is, with the appropriate background evidence, it is very probable that it is true that X is unequal, conditional on it being true that X is unfair. Likewise, with the appropriate background evidence, the fact that X is unequal will be part of the Best Explanation for why X is unfair. So, according to either of our best theories of evidential support (i.e., probabilism and explanationism; see §1), the ethical proposition that X is unfair can be good (fallible) evidence for the empirical conclusion that X is unequal or uneven. Moreover, it is plausible that one can gain epistemic justification for believing that X is unfair which does not wholly counterfactually depend on one's prior epistemic justification for believing that X is unequal (in particular, when one comes to justifiably believe that X is unfair by means of reliable testimony). Finally, and most importantly, even if we put aside all of the theoretical justifications for the epistemic permissibility of this kind of inference, it should just be intuitively clear that this inference can *often* be the right kind to give one epistemic justification for believing the conclusion. *Unfair to Unequal* gives a general schema where it is extremely plausible that ethical propositions can be legitimately used as evidence for empirical conclusions.

There are many other plausible examples. For instance, consider a child who is taught in school that littering is morally wrong, and infers from this ethical proposition to an empirical conclusion about the wrong-making features of littering (e.g., that it harms animals). This suggests the following general inference-schema:

Wrong to Wrong-Making

X is morally wrong.

Therefore, X (probably) has property Y (where Y is an empirical property which would partly explain why X is morally wrong).

Likewise, consider any case where one learns that a person is virtuous or vicious in a particular respect, and uses this piece of information to predict how this person will act. For instance, suppose I learn from a trusted and reliable news source that a particular politician is especially virtuous. I can justifiably infer from this ethical proposition to the conclusion that this politician is likely or unlikely to perform particular actions, e.g., take bribes. So, we also have the following inference-schema:

Virtue to Action

S is virtuous (in some respect).

Therefore, S is likely to act in ways which are consistent with this virtuous disposition.

These schemas are by no means exhaustive of all of the possible permissible inferences from an ethical premise to an empirical conclusion, but they should give a flavour of what these inferences should generally look like (as well as showing how they clearly do not involve any form of wishful thinking). For what it's worth, I think that it should not be very surprising that inferences like these ones are epistemically permissible. In fact, I think that we already make inferences like these ones *so routinely* that we do not even notice that we are doing it. In sum, given the plausibility of these examples of permissibly using ethical propositions as evidence for empirical conclusions, it should be clear that there are no grounds for turning the argument of the first part of this paper into a *reductio ad absurdum* of the standard metaethical assumptions which entail the Ethical Evidence thesis.

Although considering the case of an individual believer is sufficient for defusing the *reductio ad absurdum*, it is also worth outlining some examples where it is plausible that *scientists* (or the scientific community) can collectively rely upon their ethical evidence. That is, on the one hand, the potential significance of the Ethical Evidence thesis for how we should do scientific inquiry is probably the most interesting upshot of the present discussion, so it is worth exploring it in more detail. On the other hand, it might seem that it is in the specific case of scientific research that it is *least* plausible that ethical propositions can legitimately be relied upon as evidence for and against empirical conclusions, so some phantom of the *reductio* might still remain.

I do not mean to suggest that ethical propositions can be used as strong evidence for or against hypotheses in the fundamental or physical sciences. The main reason why this would be impossible is just that we are not sensitive enough to the patterns of supervenience or covariance of the truth of particular ethical propositions with the truth of propositions about fundamental physics or chemistry. However, it is far more plausible that ethical propositions can sometimes be relied upon as evidence in the *social sciences*, since we are often in a position to know (or at least justifiably believe) that the truth of particular ethical propositions systematically covaries with the truth of various propositions which feature in social scientific theories, so that learning that one of these ethical propositions is true will often make it more or less probable on our total evidence that one of these social scientific theories is true or false.

Nevertheless, it still does not appear entirely obvious that social scientists can use ethical propositions which they justifiably believe as evidence, at least, not in the same way as individuals could use these propositions as evidence in the examples above. For instance, consider a participant at a social scientific conference who, in the Q&A session on a colleague's paper on theory X, points out that even though there is robust statistical evidence in favour of theory X, the truth of this theory would be inconsistent with an ethical proposition which they take themselves to justifiably believe and therefore the theory must be false. Likewise, consider a Principal Investigator on a project who instructs their team not to bother investigating whether or not a particular phenomenon is actually occurring, since the ethical propositions which this PI justifiably believes already strongly indicate that it must be the case. It seems that the conference participant and the PI would each be doing *bad science* here. The better way for the conference participant to proceed would be to argue that there is some way in which the apparently-robust statistical evidence does not actually favour theory X after all. Likewise, the better way for the PI to proceed would be to instruct their team to find robust empirical evidence about whether this phenomenon occurs, regardless of whether or not this coheres with their ethical beliefs. So, it looks like ethical propositions cannot be legitimately used as evidence in social science after all. Moreover, we might be worried that this points towards a deeper problem with the Ethical Evidence thesis.

I have three responses to this worry. The first is the simple point that bad science need not be bad epistemology. That is, the norms of collaborative scientific inquiry sometimes

place constraints on what kinds of evidence are correct or acceptable for scientists to rely upon in particular contexts.²⁵ For instance, in most social scientific contexts, one should support their hypotheses (or challenge colleagues' hypotheses) by citing statistical data, case studies, formal models, and other stereotypical kinds of social scientific evidence. And conversely, it is often incorrect for scientists (qua scientists) to rely on evidence which is not public in the relevant sense, or at least evidence which is controversial within the scientific community. Ethical propositions will rarely fit the pattern of one of these stereotypical sources of social scientific evidence, and the truth of particular ethical propositions will often be controversial within the scientific community. Therefore, the norms of collaborative scientific inquiry will sometimes prevent scientists from correctly using ethical propositions which they justifiably believe as evidence, but this does not mean that it would be *epistemically* impermissible to use these ethical propositions as evidence in this context.

The second response is to point out that, even if there are some contexts of social scientific research where it is intuitively impermissible to rely upon ethical propositions as evidence which directly confirms or disconfirms a particular hypothesis (as in the examples above), there are many other contexts in which it is plausible that ethical propositions can legitimately play a less direct evidential role. That is, as well as playing the obvious role of directly confirming hypotheses, propositions which count as scientific evidence can also play a host of other indirect evidential roles, including directing research (e.g., by framing research questions), determining which results are controversial (and what strength of evidence is required before a hypothesis is accepted or rejected), determining how data should be collected and analysed (e.g., by indicating where we should look for relevant data, and which controls to impose on data analysis), and deciding when and where to end an analysis. Elizabeth Anderson (2004:12-18) gives a detailed case study of how value-judgements (i.e., ethical beliefs) have legitimately played all of these roles in feminist research on divorce by Abigail Stewart, Anne Copeland, Nia Lane Chester, Janet Malley and Nicole Barenbaum (Stewart et al. 1997). In particular, the idea is that starting from the justified belief in the ethical proposition *that divorce is not inherently morally bad*, feminist social scientists can direct their study

²⁵ And these norms of scientific inquiry might be pragmatically justified, e.g., by the role of science in a democracy (Bright 2018).

toward finding out the truth about the social effects of divorce (e.g., that the unhappiness of divorcees is often caused by the factors which led a divorce rather than the divorce itself, and that there are also some positive personal and social effects from divorce). So, it looks like ethical propositions can legitimately influence social scientific inquiry in just the way we would expect if they were playing a wide array of indirect evidential roles (see also, Clough 2003, 2011, Goldenberg 2015).

Finally, the third response is to outline a general kind of case where it *is* plausible that social scientists can rely upon ethical propositions which they justifiably believe as *direct* evidence for or against a social scientific hypothesis. The point is that, even if ethical evidence does not count for much in the face of more typical sources of social scientific evidence (e.g., statistical data), ethical evidence can nevertheless play a direct evidential role when these other sources of evidence run out (or alternatively, in the interpretation of these other kinds of evidence). The general kind of case I have in mind involves inferring (or refraining from inferring) causation from statistical correlation. Ethical propositions can play a direct evidential role here by counting in favour or against a causal hypothesis, when other sources of evidence run out. The specific example which I discuss below is artificially simplified in important respects, but it should do fine for the purpose of illustration.

In sociology, *Culture of Poverty* hypotheses are periodically put forward as potential explanations of why particular communities become and remain impoverished.²⁶ The basic idea here is that poor communities become and remain poor because they have bad habits and patterns of behaviour. Put quite crudely, a simple Culture of Poverty

²⁶ The Culture of Poverty hypothesis is usually taken to have been first developed by Oscar Lewis (1959). It is not entirely clear whether Lewis and other early exponents used the theoretical term, 'culture', to refer to a broad social structural phenomenon, or to individual members of a community's habits and values (Steinberg 2011). The latter became the standard understanding of the hypothesis, especially after the publication of the notorious Moynihan Report (1965), which argued that one important reason why African-American communities remained poor was the fact that their 'family values' differed from those of White Americans, in such a way that prevented African-Americans from escaping poverty. This met with immediate criticism and controversy but shaped popular discourse about the 'culture of poverty' with long-lasting effects. The theory re-emerged in the 1980's under the term, 'underclass' (Auletta 1982), and it has recently sprung up again (e.g., Small et al., 2010); though in a much more nuanced form than I do justice to in the following discussion. Additionally, individualistic versions of the Culture of Poverty hypothesis appear to remain prevalent in certain circles of popular right-wing discourse, especially in the USA – often with clearly racist undertones.

hypothesis says something like: actual minority ethnic communities have a disproportionate tendency to become and remain poor *because* members of these communities are disproportionately likely not to work hard, to have a culture of depending on welfare payments, to spend money on luxuries rather than investing, and so on. There are also more sophisticated versions of the Culture of Poverty hypothesis, which are not quite as individualistic as this crude version. In order to keep things simple, however, I focus on the crude version below.

Although there may be other reasons to think that the simplistic version of the Culture of Poverty hypothesis is false, it turns out that it is not entirely straightforward to refute empirically.²⁷ That is, there may be robust statistical data which suggest that many members of poor communities *do* have different attitudes towards work and patterns of spending behaviour from members of relatively wealthy communities. The challenge is to evaluate whether these patterns of behaviour are the *cause* of poverty (and the Culture of Poverty theory is true) or merely its *symptom* (and there is a more structural, less individualistic, explanation for why certain communities become and remain poor).

As I outline below, one good reason to conclude that the Culture of Poverty hypothesis is false is that this kind of individualistic explanation of poverty is in tension with an ethical proposition which we justifiably believe, namely, *that particular actual patterns of wealth inequality in highly developed countries are unjust*. However, I first want to stress that this does *not* mean that it would be impossible to find straightforward empirical evidence against the Culture of Poverty hypothesis (although this might take up a lot of work and resources). What the ethical evidence offers us is not necessarily *better* evidence against the sociological hypothesis, but rather evidence which we happen to *already possess*. In other words, when the typical sources of social scientific evidence run out, we still possess some other resources for deciding between different hypotheses, and these

²⁷ This is probably an exaggeration. However, the relevant point is just that *even if* our current sociological evidence could not definitively reject the Culture of Poverty hypothesis, it would still be plausible that we could use our justified belief in an ethical proposition as further evidence against this hypothesis.

resources include the ethical propositions which we justifiably believe and therefore possess as good evidence for empirical conclusions.²⁸

At the risk of belabouring this example, here are the two relevant propositions, namely, a statement of a simple version of the Culture of Poverty hypothesis (CoP) and a particular version of the ethical statement *that particular patterns of inequality are unjust* (E):

CoP: Poor communities typically become/remain poor *because* individual members of these communities are disproportionately likely to have a poor work ethic, unwise spending habits, a culture of welfare-dependence, etc.

E: The actual, contingent state of affairs where some (majority) non-White communities in the USA are persistently poorer and have less social mobility than (majority) White communities, is unjust.

Let's assume that we, as scientists, have some epistemic justification for believing CoP, for instance, that there is a documented empirical correlation between persistent poverty and these individual character traits. We also (as a matter of fact) have clear justification for believing E, through an obscure but reliable process of diffuse testimony within our ethical community, which plausibly partially originates in the lived experience of those who have experienced this unjust poverty. It should be clear that there is *at least some* tension between these two propositions. As I spell out below, taking CoP to be true should lead us to lower our confidence that E is true, and taking E to be true should lead us to correspondingly lower our confidence in CoP.

The first thing to note is that CoP and E are not strictly speaking inconsistent; it might be possible for the actual racial inequality in the USA (and elsewhere) to be unjust even though poverty is generally caused and sustained by individual character traits. However, upon making some plausible background assumptions about the connection between injustice, responsibility and motivation, CoP and E *do* become mutually inconsistent. In particular:

²⁸ Another way of making this point is in terms of a Quinean 'web of belief', where the total body of possessed evidence against which we epistemically evaluate our total body of theories includes our justified ethical beliefs (see especially White 1981, Nelson 1990, Solomon 2012).

BE: If actual racial inequality is unjust, then the best explanation currently available to us for why this is the case includes (or entails) the fact that members of some poor communities are not responsible for their own persistent poverty.

R: If CoP gave a true description of how a particular community became/remains poor, then individual members of this community would be (at least partly) responsible for their own persistent poverty.

In other words, if E and BE are true, then we have justification for believing that members of some actual poor communities are not responsible for their own poverty. If CoP and R are true, then we have justification for believing that individual members of these communities *are* responsible for their own poverty. We have a contradiction here, so we must reject at least one of the propositions E, CoP, BE or R. Or, more realistically, we must decrease our degree of confidence in one or more of these propositions.

Although BE and R could certainly be challenged (and such challenges deserve to be taken seriously), I hope that my reader sees why they are at least *prima facie* plausible. BE says that the most plausible explanation for why actual inequality is unjust involves the claim that the structural features which cause poverty are out of the control of those who are poor. R points out that we can understand simple versions of the Culture of Poverty hypotheses as implying that an individual's persistent poverty is their own fault. For the purpose of the present discussion, therefore, I will take these two assumptions for granted.

The point is that, *insofar as* we (collectively, as scientists) are confident that E and CoP are mutually inconsistent, we should (collectively) correspondingly decrease our confidence that one or the other of these propositions is true. Which belief should we keep and which should we revise? I do not want to suggest that there is necessarily only one rational answer here. Rather, I am suggesting that it can *sometimes* be rational to decrease our confidence in CoP rather than decreasing our confidence in E. In other words, it would *not* be epistemically irrational to rely upon our justified belief in E in order to reject (or decrease our degree of confidence in) CoP. Not only is this ethical proposition the right sort of thing *metaphysically* to count as evidence against CoP (§1), but our collective epistemic justification for believing this proposition is also at least

partly independent of our antecedent justification for believing the relevant empirical propositions (because this ethical belief is formed via a diffuse testimonial process), so E is also *epistemically* the right sort of thing to be used as evidence against CoP in this particular context (§2). Therefore, it *can* be rational for us to use the ethical proposition *that actual patterns of inequality are unjust* as evidence against a social scientific hypothesis which posits an individualistic explanation of persistent poverty (and conversely, as evidence in favour of a structural explanation of this poverty).

4. Conclusion

In the philosophical debate about the role of value-judgements in scientific inquiry, it is a dogma that ethical propositions cannot play a straightforward evidential role in scientific research. Genealogically, this dogma is probably rooted in the extreme forms of metaethical non-cognitivism which were popular amongst analytic philosophers in the mid- to late-20th century (c.f., Brown 2020:101-103). These metaethical theories are now rightly seen to be highly implausible, and in the last 20 years we have seen a move towards realism in the metaethical debate (or, at least, towards accommodating realist intuitions about ethical thought and talk). As a result (as I argued in §1), there are no grounds for thinking that the metaphysical weirdness of ethical propositions prevents them from counting as evidence for or against empirical conclusions. Moreover (as I argued in §2), when we think about ethical epistemology as a form of social epistemology, where many of our ethical beliefs are justified at least partly via testimony, then these beliefs will often be sufficiently independently justified for us to be in a position to rationally rely upon the contents of these beliefs as evidence for and against empirical conclusions. Finally (as I outlined in §3), this should not be seen as an absurd result. There are plenty of cases where it is intuitively epistemically permissible for an individual to use an ethical proposition as empirical evidence, and moreover, it is also plausible that there are some contexts of (social) scientific inquiry where ethical propositions can play both direct and indirect evidential roles.

In conclusion, then, even though the metaphysics and epistemology of the ethical might be significantly weirder, in some respects, than the metaphysics and epistemology of the

empirical, it is just not weird enough to prevent ethical propositions from being apt to be relied upon as evidence for empirical conclusions. And we can (and should) exploit this ethical evidence in order to gain a better understanding of important aspects of empirical reality.²⁹

References:

1. Anderson, Elizabeth (2004). Uses of value judgments in science: A general argument, with lessons from a case study of feminist research on divorce. *Hypatia* 19 (1):1-24. [10.1111/j.1527-2001.2004.tb01266.x](https://doi.org/10.1111/j.1527-2001.2004.tb01266.x)
2. Anderson, Elizabeth (2016). The Social Epistemology of Morality. In Michael Brady & Miranda Fricker (eds.) *The Epistemic Lives of Groups: Essays in the Epistemology of Collectives*. Oxford: Oxford University Press. pp.75-94
3. Auletta, Ken (1982). *The underclass* (1st ed.). New York: Random House.
4. Barber, Alex (2013). Science's Immunity to Moral Refutation. *Australasian Journal of Philosophy* 91 (4):633-653. [10.1080/00048402.2013.768279](https://doi.org/10.1080/00048402.2013.768279)
5. Basl, John & Coons, Christian (2017). Ought to Is: The Puzzle of Moral Science. *Oxford Studies in Metaethics* 12.
6. Bedke, Matthew (2017). Cognitivism and Non-Cognitivism. In Tristram McPherson & David Plunkett (eds.), *The Routledge Handbook of Metaethics*. Routledge. pp. 292-307.
7. Bright, Liam (2018). Du Bois' democratic defence of the value free ideal. *Synthese* 195 (5):2227-2245.
8. Brown, Jessica (2016). Contextualism about Evidential Support. *Philosophy and Phenomenological Research* 92 (2):329-354. [10.1111/phpr.12164](https://doi.org/10.1111/phpr.12164)
9. Brown, Matthew J. (2013). Values in Science beyond Underdetermination and Inductive Risk. *Philosophy of Science* 80 (5):829-839. [10.1086/673720](https://doi.org/10.1086/673720)
10. Brown, Matthew J. (2020). *Science and Moral Imagination: A New Ideal for Values in Science*. Pittsburgh: University of Pittsburgh Press.
11. Campbell, Richmond (1998). *Illusions of Paradox: A Feminist Epistemology Naturalized*. Rowman & Littlefield Publishers.
12. Clough, Sharyn (2003). *Beyond Epistemology: A Pragmatist Approach to Feminist Science Studies*. Lanham, Md.: Rowman & Littlefield Publishers.

²⁹ Many thanks to Al Prescott-Couch for extensive discussion around this project and suggestions for some of the examples. Thanks also to Jonathan Ichikawa, Emily Tilton, Phyllis Pearson, Alex Bryant, Kelsey Vicars, Cam Gilbert, Mira Kuroyedov and two reviewers for this journal.

13. Clough, Sharyn (2011). Radical Interpretation, Feminism, and Science. In Jeffrey Malpas (ed.) *Dialogues with Davidson*, pp.405–426. Cambridge: MIT Press
14. Crisp, Roger (2014). II—Moral Testimony Pessimism: A Defence. *Aristotelian Society Supplementary Volume* 88 (1):129-143. [10.1111/supa.2014.88.issue-1](https://doi.org/10.1111/supa.2014.88.issue-1)
15. Dorr, Cian (2002). Non-cognitivism and Wishful Thinking. *Noûs* 36 (1):97–103. [10.1111/1468-0068.00362](https://doi.org/10.1111/1468-0068.00362)
16. Douglas, Heather (2000). Inductive risk and values in science. *Philosophy of Science* 67 (4):559-579. [10.1086/392855](https://doi.org/10.1086/392855)
17. Douglas, Heather (2009). *Science, Policy, and the Value-Free Ideal*. Pittsburgh: University of Pittsburgh Press.
18. Elliott, Kevin C. (2017). *A Tapestry of Values: An Introduction to Values in Science*. Oxford University Press USA.
19. Geertz, Clifford (1990). A lab of one's own. *New York Review of Books* 37 (8 November): pp.19-24
20. Gibbard, Allan (2003). *Thinking How to Live*. Harvard University Press.
21. Goldenberg, Maya J. (2015). How Can Feminist Theories of Evidence Assist Clinical Reasoning and Decision-Making? *Social Epistemology* 29 (1): 3–30. [10.1080/02691728.2013.794871](https://doi.org/10.1080/02691728.2013.794871)
22. Groll, Daniel & Decker, Jason (2014). Moral Testimony: One of These Things Is Just Like the Others. *Analytic Philosophy* 55 (1):54-74. [10.1111/phib.12033](https://doi.org/10.1111/phib.12033)
23. Gross, Paul, and Norman Levitt (1994). *Higher superstition: The academic left and its quarrels with science*. Baltimore: Johns Hopkins.
24. Haack, Susan (1993). Epistemological reflections of an old feminist. *Reason Papers* 18 (fall) pp.31-43
25. Hills, Alison (2009). Moral testimony and moral epistemology. *Ethics* 120 (1):94-127. [10.1086/648610](https://doi.org/10.1086/648610)
26. Hills, Alison (2013). Moral Testimony. *Philosophy Compass* 8 (6):552-559. [10.1111/phc3.12040](https://doi.org/10.1111/phc3.12040)
27. Lackey, Jennifer (2006). It takes two to tango: beyond reductionism and non-reductionism in the epistemology of testimony. In Jennifer Lackey & Ernest Sosa (eds.), *The Epistemology of Testimony*. Oxford University Press. pp. 160--89.
28. Lewis, Oscar (1959). *Five families: Mexican case studies in the culture of poverty*. Basic Books.
29. Lipton, Peter (2003). *Inference to the Best Explanation*. Routledge.
30. Longino, Helen (1990). *Science as social knowledge : Values and objectivity in scientific inquiry*. Princeton: Princeton University Press.
31. Majors, Brad (2007). Moral explanation. *Philosophy Compass* 2 (1):1–15. [10.1111/j.1747-9991.2006.00049.x](https://doi.org/10.1111/j.1747-9991.2006.00049.x)
32. McCain, Kevin (2013). Explanationist evidentialism. *Episteme* 10 (3):299-315. [10.1017/epi.2013.22](https://doi.org/10.1017/epi.2013.22)
33. McPherson, Tristram (2012). Ethical Non-Naturalism and the Metaphysics of Supervenience. In *Oxford Studies in Metaethics Vol 7*. pp. 205.

34. McShane, Paddy Jane (2021). Against epistemic pessimism about moral testimony. *Episteme* 18 (2):200-223. [10.1017/epi.2019.11](https://doi.org/10.1017/epi.2019.11)
35. Mogensen, Andreas L. (2017). Moral Testimony Pessimism and the Uncertain Value of Authenticity. *Philosophy and Phenomenological Research* 95 (2):261-284. [10.1111/phpr.12255](https://doi.org/10.1111/phpr.12255)
36. Moretti, Luca & Piazza, Tommaso (2013). When warrant transmits and when it doesn't: towards a general framework. *Synthese* 190 (13):2481-2503. [10.1007/s11229-011-0018-2](https://doi.org/10.1007/s11229-011-0018-2)
37. Moynihan, Daniel Patrick. (1965). The Negro family: The case for national action. *Washington, DC: Office of Policy Planning and Research, U.S. Department of Labor.*
38. Nelson, Lynn Hankinson (1990). *Who Knows: From Quine to a Feminist Empiricism*. Temple University Press.
39. Podgorski, Abelard (2020). Normative Uncertainty and the Dependence Problem. *Mind* 129 (513):43-70. [10.1093/mind/fzz048](https://doi.org/10.1093/mind/fzz048)
40. Rosen, Gideon (2017). Metaphysical Relations in Metaethics. In Tristram McPherson & David Plunkett (eds.), *The Routledge Handbook of Metaethics*. Routledge. pp. 151-169.
41. Sliwa, Paulina (2012). In defense of moral testimony. *Philosophical Studies* 158 (2):175-195. [10.1007/s11098-012-9887-6](https://doi.org/10.1007/s11098-012-9887-6)
42. Small, Mario Luis, Davis Harding, and Michele Lamont, eds. (2010). Reconsidering culture and poverty. *ANNALS of the American Academy of Political and Social Science* 629.1
43. Solomon, Miriam (2012). The Web of Valief : An Assessment of Feminist Radical Empiricism. In (Ed.), *Out from the Shadows: Analytical Feminist Contributions to Traditional Philosophy.* : Oxford University Press
44. Steinberg, Stephen (2011, January 13). Poor Reason: Culture still doesn't explain poverty. *Boston Review*. Retrieved from <https://bostonreview.net/articles/steinberg-php/> 11 April 2022
45. Stewart, Abigail, Anne Copeland, Nia Chester, Janet Malley, and Nicole Barenbaum (1997). *Separating together: How divorce transforms families*. New York: Guilford Press.
46. Toole, Briana (2021). Recent Work in Standpoint Epistemology. *Analysis* 81 (2):338-350. [10.1093/analys/anab026](https://doi.org/10.1093/analys/anab026)
47. Tucker, Chris (2010). When Transmission Fails. *Philosophical Review* 119 (4):497-529. [10.1215/00318108-2010-012](https://doi.org/10.1215/00318108-2010-012)
48. Werner, Preston J. (2020). Moral perception. *Philosophy Compass* 15 (1). [10.1111/phc3.12640](https://doi.org/10.1111/phc3.12640)
49. White, Morton Gabriel (1981). *What is and What Ought to Be Done: An Essay on Ethics and Epistemology*. Oxford University Press.
50. Wiland, Eric (2021). *Guided by Voices: Moral Testimony, Advice, and Forging a 'We'*. Oxford, UK: Oxford University Press.

51. Williamson, Timothy (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
52. Wright, Crispin (2002). (Anti-)sceptics simple and subtle: G. E. Moore and John McDowell. *Philosophy and Phenomenological Research* 65 (2):330-348.
[10.1111/j.1933-1592.2002.tb00205.x](https://doi.org/10.1111/j.1933-1592.2002.tb00205.x)