

A First-Order Argument for Evidence-Constrained Morality

Alex Worsnip

1. Two cases

Let's begin with two cases:

Light Switch.¹ You walk into your office on a Monday morning, and flip the light switch. Flipping the switch causes the building's electrical system to short-circuit and explode, causing the whole building to go up in flames, killing everyone inside. You had no way of anticipating that this would happen.

Anti-Vaxxer. It is late 2020, at the peak of the COVID-19 pandemic. You believe, contrary to your evidence, that the pandemic is a hoax, and that the vaccine is designed to implant microchips into people so that Bill Gates can track their movements. You advise all your friends and family, including your elderly and immunocompromised relatives, to go about their daily business without masks, and to refuse the vaccine, and they follow your advice. Over the subsequent months, several of your elderly relatives contract COVID-19 and die.

My judgments about these cases, which I hereby invite you to share, are these: in flipping the switch in **Light Switch**, you do not act (morally) wrongly. But in advising your relatives as you do in **Anti-Vaxxer**, you do act (morally) wrongly.

2. Three (very rough) theories

With that on the table, let's consider two views about moral obligation² that are the subject of a classic debate:

Objectivism (a.k.a. the "fact-relative" view). Our moral obligations are determined by the full set of (non-normative) facts, regardless of our epistemic access to them.³

Subjectivism (a.k.a. the "belief-relative" view). Our moral obligations are determined by our beliefs about the (non-normative) facts, regardless of whether they are justified by our evidence.⁴

I don't love the labels 'objectivism' and 'subjectivism'. First, these terms are used in many other ways, and are thus liable to confuse. The debate at issue here has little to do with the controversy between objectivism and subjectivism in the sense connected with metaethical realism and anti-realism, for

Thanks to audiences at Syracuse University, a workshop on metaethics at Harvard, a workshop on the new ethics of belief at KU Leuven, and the Law & Philosophy Speaker Series at Yale Law School. Special thanks to Daniel Fogal, Zoë Johnson King, Raúl Ibarra Herrera, Maggie Shea, and multiple anonymous referees.

¹ This case is based on one from Thomson (1990: 229).

² I'll assume throughout that acts contrary to one's moral obligations are morally wrong, and that acts that are not contrary to one's moral obligations are morally permissible. Hence, though I state the views in terms of obligation, they have immediate consequences for other deontic notions such as wrongness and permissibility.

³ See e.g. Moore (1903), Thomson (1986, 1990), and Graham (2010, 2021).

⁴ See e.g. Ross (1939) and Prichard (1949/2002).

example. Second, this terminology also suggests a bipartite division between possible views, which (as we'll see in a moment) isn't right. I prefer the 'fact-relative' and 'belief-relative' labels.

What's meant by "determined by" as it occurs in the statements of these views? The answer is this: notwithstanding the myriad metaethical controversies about grounding, supervenience, naturalism vs. non-naturalism, and so on, pretty much everyone agrees that normative facts about what's required of you in *specific situations* (the "particular" normative facts) at least partly depend on certain normatively relevant but (themselves) non-normative facts. 'Determined by' is just a stand-in for whatever metaphysical relation obtains between those non-normative facts and the normative facts. The controversy is about *which* non-normative facts—or, if we go with the belief-relative view, perhaps which propositions—are eligible to stand in this metaphysical relation with a particular class of normative facts, viz. those about moral obligation.

With that said: prima facie, **Light Switch** is a counterexample to objectivism, since it is a case where a certain fact you lack epistemic access to (viz. that flipping the switch will blow up the building) is intuitively not eligible to determine whether you violate an obligation in acting (by flipping the switch). Conversely, **Anti-Vaxxer** is a counterexample to subjectivism, since it is a case where *unjustified* beliefs of yours (viz. that the pandemic is a hoax, that the vaccine is designed to implant microchips into people so that Bill Gates can track their movements) are not eligible to make (or contribute toward making) your action of advising your family as you do permissible—even though, presumably, there's a sense in which those beliefs if true or justified *would* make it right to advise them as such.

The pair of cases, however, suggests an intermediate view.

(Evidential) Perspectivism (a.k.a. the "evidence-relative" view). Our moral obligations are determined by our evidence about the non-normative facts.⁵

Light Switch and **Anti-Vaxxer** together suggest that this is the "Goldilocks" view. A weird feature of early literature on this topic is that it doesn't consistently distinguish the belief-relative and evidence-relative views, with either potentially going under the label "subjectivism".⁶ Cases like Anti-Vaxxer show they are clearly (importantly) different, hence yielding an (at least) tripartite (not bipartite) division between views. And they also suggest that the evidence-relative view is prima facie superior to the belief-relative one.

The evidence-relative view is also supported by analogies to the received views about (substantive) rationality.⁷ First, on the *epistemic* side of substantive rationality, the vast majority of epistemologists hold that what it's (substantively) rational to believe is determined by your evidence—not by all the facts (as would be the case if you were rationally required to believe all the truths and

⁵ See e.g. Zimmerman (2006, 2008, 2014); Scanlon (2008: 50-1); Kiesewetter (2011). Different formulations of and variations on the view are possible: for example, it might be said that our moral obligations are determined by the subset of the non-normative facts that are part of our evidence, or by the subset of the non-normative facts that are supported by our evidence. The subtle differences between the view as I've stated it in the text and these other views won't be important for my purposes.

⁶ Feldman (1988) has a nice catalogue of the instances in the literature where "subjectivism" refers to the evidence-relative view and those where it refers to the belief-relative view.

⁷ Substantive rationality concerns one's responsiveness to reasons. It contrasts with structural rationality, which is concerned with coherence between one's mental states. See e.g. Fogal & Worsnip (2021).

disbelieve all the falsehoods), nor by your beliefs, no matter how poorly evidentially supported those beliefs (as would be the case if groundlessly believing p sufficed to make it *substantively* rational to believe q when p entails q). Second, on the *practical* side of substantive rationality, there is increasing consensus that substantive rationality is a matter of acting in accordance with your evidence-relative reasons, not your fact-relative reasons or your belief-relative reasons.⁸ It isn't *irrational* to fail to respond to reasons you couldn't possibly have known of: if the glass in front of you contains petrol but all your evidence suggests it contains gin and tonic, it's not irrational to drink. But equally, it doesn't suffice for the rationality of drinking that you merely *believe* that it contains gin and tonic if all your evidence suggests it contains petrol.

Evidence-relative views about (substantive) rationality effectively posit a constraint on the kinds of facts or propositions that can determine what is rational for you: they must be part of, or supported by, your evidence. And very plausibly this is a *conceptual* constraint on the notion of rationality: while there might be some sense in which you *ought* not to drink the glass of petrol (even if all your evidence suggests it's gin and tonic), it would be a conceptual error to say that this 'ought' is one of rationality, and hence that it would be irrational to drink.

So too the evidential perspectivist view posits a (perhaps conceptual) constraint on the notions of moral obligation and moral wrongness: while there may be some sense in which you ought not to flip the switch in **Light Switch**, it would be a (perhaps conceptual) error to say that this 'ought' is a moral one—or, at the very least, it would be an error to say that it expresses the concept of moral *obligation*—and hence, it would likewise be an error to say that it would be *wrong* to flip the switch.⁹

One final clarification about evidential perspectivism. As I understand it, the debate between objectivists and perspectivists is about whether what morality *primarily, non-derivatively* demands of us is constrained (in a certain way) by our evidence. Evidential perspectivism is the claim that it is. Thus, my perspectivist does not think of evidence-relative moral demands as explanatorily derivative on some set of more objective, evidentially-unconstrained moral demands.¹⁰ That is, the “*realest*”, *most objective moral demands there are* are themselves evidentially constrained.^{11, 12}

⁸ Cf. Kiesewetter (2017), Lord (2018), and Fogal & Worsnip (2021).

⁹ What kind of 'ought' is it if not a moral 'ought'? My suggestion is that it's an evaluative 'ought'—one indexed to what would make things go best. The existence of such evaluative 'ought's is well-recognized (e.g. Williams 1981; Wedgwood 2007: 117; Broome 2013: 18); they are the kind of 'ought' at issue in (for example) 'there ought not to be so much suffering in the world'. It's distinctive of evaluative 'ought's that they need not concern what any particular agent ought to do, though we *can* apply them to the evaluation of agents' acts.

¹⁰ Contra e.g. Ewing (1947) and Smith (2018). Some perspectivists may concede this explanatory derivativeness but hold that it's the 'ought' expressed by the evidence-relative moral demands that is the “deliberatively relevant” one (and that it's in virtue of holding this that their view is perspectivist). Cf. Jackson (1991: 472) and Kiesewetter (2011: 2). But see Pittard (2023) for some skepticism about a single, privileged deliberatively relevant 'ought'.

¹¹ Again, this is consistent with the view that 'ought' can be used in an evidentially-unconstrained way; it just denies that this is a *moral* ought. It is also consistent with the view that there are further moral 'ought's that are derivative on the (primary) *evidentially-constrained* one—perhaps ones that are *belief*-relative, or that take account of time or reasoning limitations. For both these reasons, the view defended here is consistent with pluralism and/or contextualism about 'ought' (Björnsson & Finlay 2010; Dowell 2013; Pittard & Worsnip 2017).

¹² This can make it seem hard to characterize the formal difference between fact-relative and evidence-relative views. On any view, some facts are morally relevant while others aren't: no-one thinks it's literally the case that *all* facts are morally relevant to what whether I ought to Φ , for all Φ . For example, whether it's currently raining in Timbuktu isn't relevant to whether I should keep a promise to my daughter to buy her a bike for Christmas. So, one might think, the fact-relative view is just that our moral obligations are determined by all the *morally relevant* facts. But on one way of thinking about the

As so often, utilitarianism—not because it is plausible but because it is helpfully concrete—provides a handy illustration. Some utilitarians think that you morally ought to do whatever (actually) maximizes *actual* utility. Other utilitarians think that you morally ought to do what maximizes *expected* utility, where expected utility is a function not of the outcomes that various options *will* produce but rather of your *evidence* about what outcomes they will produce, or slightly more precisely, the evidential probabilities of each outcome conditional on a certain action’s being performed. The former is a fact-relative version of utilitarianism, while the latter is an evidence-relative version. Now: does an advocate of the view that you ought to maximize *expected* utility (given your evidence) need to say that this ‘ought’ is derivative on a more “primary” ‘ought’ requiring you to maximize actual utility? No! It can instead be explained by facts about objective *goodness* or *value*.¹³ On this view, maximizing actual utility is what is *valuable*, yet since the moral ‘ought’ is evidentially constrained (in a way that facts about value are not), what you morally ought to do—in the most primary, objective, “real” sense there is—is maximize *expected* utility. And this is so precisely because this is the most reliable, feasible way to achieve what is of value, given your evidence.

3. Problems, problems

My judgments about cases like **Light Switch** and **Anti-Vaxxer**—and other analogous ones—are very firm, and so far it seems (at least to me) that the evidence-relative view is the clear winner—as I said, the “Goldilocks” view. But there are several known problems for evidential perspectivism. Here I’ll focus on four (in what seems to me roughly ascending order of importance).¹⁴

First, not everyone shares these intuitions. Some (defending the fact-relative view) say that you act *wrongly* in **Light Switch**, but are not *blameworthy*.¹⁵ Some (defending the belief-relative view) might say that in **Anti-Vaxxer**, you’re criticizable for your beliefs, but not your resulting actions.

Second, some of the intuitions that drive rejection of the fact-relative view may be exploitable in other cases to pressure the evidence-relative view. For example, some common complaints against the fact-relative view are that it isn’t “followable” or doesn’t provide “guidance,”¹⁶ or that one can violate its prescriptions non-criticizably (as in **Light Switch**). But there are cases where this also seems to be true of the evidence-relative view. For example, suppose that what your evidence supports doing on balance is a complicated matter that takes time to figure out (perhaps it involves complex expected utility calculations). But suppose you have no time: the building is burning, and you need to decide whether to rush in. Here arguably the evidence-relative view also isn’t followable, and one might

evidence-relative view, it just holds that facts outside an agent’s epistemic grasp are *normatively irrelevant* to what she should do much as the fact about whether it’s raining in Timbuktu isn’t relevant to whether I should keep my promise. And in that case, it’s consistent with the letter of the view that our moral obligations are determined by all the *morally relevant* facts. Still, though, the objectivist and (this) perspectivist disagree on whether some fact’s being outside an agent’s epistemic grasp automatically makes it irrelevant to what they should do. That’s enough for a meaningful disagreement.

¹³ Cf. Zimmerman (2006: 341); Kiesewetter (2011: 19-20).

¹⁴ One problem I won’t discuss is the problem of accounting for advice from a superior epistemic position (cf. Thomson (1986: 184; Graham 2010: 91-92; Kolodny & MacFarlane 2010). Others have responded to that problem ably (Zimmerman 2006: esp. 345 n. 40; Gibbons 2010; and especially Kiesewetter 2011). My own response would largely mirror theirs, and not turn significantly on my particular version of perspectivism or arguments for it.

¹⁵ E.g., Moore (1903), Thomson (1990), and Graham (2010: 94).

¹⁶ See, e.g., Scanlon (2008: 48); Smith (2010: 67-68, 2018); Kiesewetter (2011: 5-6); Lord (2015); Fox (2019). Cf. also Prichard (1949/2002: 87-90).

violate its prescriptions non-criticizably.¹⁷ In fact, some argue that *no* interesting norm is perfectly followable or impossible to violate non-criticizably.¹⁸

Third, the evidential perspectivist may seem to have trouble handling cases where you ought to (have) put yourself in a better epistemic position.¹⁹ Such cases are ones where given your actual evidence, Φ -ing is the thing to do, but you ought to have evidence that would have made something else the thing to do. For example, consider a modified version of **Anti-Vaxxer** where you simply have *ignored* the evidence that demonstrates the reality of COVID, the safety and efficacy of vaccines, etc., and so don't possess this evidence (while you have gathered much of the evidence that tells in the other direction). There are two issues here. First, we want to explain why you have reason to put yourself in a better epistemic position in such a case. The worry here is that if your moral obligations are relative to your existing evidence, you never need to gather more evidence to know what your moral obligations are. Second, we may want to account for the thought that if you culpably *fail* to gather more evidence in such cases, your eventual *act* (here, advising your family not to get vaccinated) is wrong. The worry here is that since the evidence you *do* have favors this act, evidential perspectivism can't explain this.

Fourth and finally—and this is perhaps the deepest problem—I have formulated evidential perspectivism (and its rivals) as a view about which *non-normative* facts (help to) determine our moral obligations. (For example: whether flipping the switch will cause an explosion, whether COVID-19 is a hoax, whether the vaccine implants microchips, and so on.) But nevertheless, we can ask the following question: suppose you have justified but false beliefs about the (“pure”²⁰) *moral facts themselves*, and act accordingly. Do you act wrongly?

On one hand, there is some pressure for the evidential perspectivist to say ‘no’. Justified false beliefs about morality are possible,²¹ just as justified beliefs about morally relevant non-moral matters are. And if justified false beliefs about morally relevant non-moral matters can get you off the hook for the charge of acting wrongly, why not justified false beliefs about morality itself? (What's more morally relevant than the moral facts?)

On the other hand, there is some pressure for the evidential perspectivist (or anyone) to say ‘yes’. The moral facts are the moral *facts*: they are by hypothesis the facts about what you *really* morally ought to do. And so it is hard to see how it could not be true that at least in *some* good sense you (really!) morally ought to do as they instruct you to, whether you have justified beliefs to the contrary

¹⁷ Cf., e.g., Kagan (2018); Lasonen-Aarnio (2024).

¹⁸ Cf., e.g., Hawthorne & Srinivasan (2013); Srinivasan (2015).

¹⁹ Cf. Smith (2014).

²⁰ By ‘pure’ moral facts, I mean the moral facts that do not themselves depend on contingent empirical facts. Cf. McGrath (2009). In general, in this paper, I talk about our evidence and justification for morally relevant non-moral claims, and about our evidence and justification for pure moral claims. It might be wondered where *impure* moral claims are in all this. But mistakes about the impure moral facts are traceable either to (perhaps implicit) mistakes about the pure moral facts or to mistakes about the morally relevant non-normative facts (or both). How mistakes about impure moral facts are to be handled, on my view, will depend on what the mistake is traceable to. Note, though, that for an evidential perspectivist, it won't be possible for a mistake about an impure moral fact to be traceable to a *justified* mistake about a morally-relevant non-moral fact. This is because the impure moral facts are fixed by one's evidence about the morally-relevant non-moral facts, not by those facts themselves. Hence, any *justified* mistake about an impure moral fact must be traceable to a justified mistake about the pure moral facts. Thanks to Daniel Fogal for pushing me to clarify all this.

²¹ Pace Smithies (2023).

or not. When it comes to the moral facts, if there is some sense in which you (morally) ought to do what you justifiably believe you morally ought to do, it seems it must be derivative on the sense in which you (morally) ought to do what you actually morally ought to do.

4. The moral obligation to be epistemically responsible

Existing arguments for perspectivism often rely on abstract, quasi-metaethical considerations about followability, guidance, or excusability.²² These arguments walk into the second problem above, which concerns precisely the instability of arguments from followability, guidance, and excusability. By contrast, as my title promises, my argument will be first-order. That is: it argues for first-order moral principles that are evidentially-constrained, on first-order grounds. I'll try to show, all in good time, that the particular view I defend solves the other problems for evidential perspectivism, too.

To begin, I'm going to apparently change the subject a bit, and articulate a view in the ethics of belief that I'll call *moralistic evidentialism*.

Moralistic Evidentialism.

- i. Being epistemically responsible involves one's beliefs being properly responsive to one's evidence.
- ii. At least when it comes to (potentially) morally relevant matters,²³ we have a moral obligation to be epistemically responsible, as such.²⁴

As is often noted, W.K. Clifford—often credited with being the first philosopher to explicitly endorse evidentialism—considers the requirement to believe in accordance with our evidence to be a moral one.²⁵ My statement of moralistic evidentialism can be thought of as a qualified and slightly more complicated version of Clifford's. The qualification is that it only applies to subject-matters that are potentially morally relevant. While this is a qualification, it nevertheless may be that the class of potentially morally relevant matters is quite large. Many facts are morally relevant (especially on consequentialist views that make all the consequences of our actions relevant to their moral status), and moreover, we often don't know in advance (when forming beliefs) whether some fact may turn out to be morally relevant to something or not.

The extra complication that moralistic evidentialism adds to Clifford's view is that it affirms that there *is* a distinctively epistemic requirement to believe in accordance with our evidence. It just adds that because we are (often) morally required to be epistemically responsible as such, we are also

²² See n. 16 above. Zimmerman (2006) also gives an argument for the evidence-relative view that doesn't appeal to considerations of followability or guidance. But even Zimmerman's argument relies on an abstract, general connection between obligation and conscientiousness that I won't have to assume.

²³ Do the morally relevant matters here include morality itself, or just the morally-relevant non-moral matters? For now, I only want to assume that they include the morally-relevant non-moral matters. But there's a good case that they should include the morality itself too. As I'll come back to in §7, this is an important choice-point for the theory I'm advocating. But my basic argument will only require assuming that we're morally required to believe what our evidence supports about morally-relevant non-moral matters.

²⁴ Note: when I use 'responsible' here, I mean it in the sense of 'please drink responsibly', not in the sense of 'you are responsible for your own actions'. In other words, 'responsible' here means that you've conducted yourself in a responsible manner, not that you are eligible for being praised or blamed.

²⁵ Clifford (1877/1999).

morally required to believe in accordance with our evidence. The ‘as such’ indicates that it isn’t just that things happen to line up such that the thing that is epistemically required is morally required; rather, whatever it is that is epistemically required (with respect to morally relevant matters) is morally required, partly in virtue of its being epistemically required.

Notice that this view is interestingly a kind of inverse of the phenomenon involved in (alleged) moral encroachment.²⁶ Roughly, moral encroachment occurs when (if ever) the moral status of our beliefs encroaches on their epistemic status – they are epistemically unjustified partly in virtue of their moral wrongness. According to moralistic evidentialism, by contrast, the epistemic status of our (morally relevant) beliefs encroaches on their moral status – they are morally wrong partly in virtue of being epistemically unjustified.

What’s the rationale for moralistic evidentialism? There are two, related, lines of thought. First, epistemically unjustified beliefs affect our actions, sometimes in disastrous ways—as in **Anti-Vaxxer**. Now, of course, they don’t always have disastrous upshots. Unjustified beliefs can be (luckily) true. Or, in other cases, they may be false but not lead to any disaster: for example, there’s a version of **Anti-Vaxxer** where your relatives don’t follow your advice, or do follow it but are lucky enough not to get badly sick. But the idea is that at least when it comes to epistemically unjustified beliefs that are morally relevant, holding them (given their potential consequences for action) is morally *negligent* or *reckless*, even when the bad consequences don’t eventuate.

As the distinction between negligence and recklessness is standardly drawn in law and philosophy, recklessness involves being aware of a risk but disregarding it unreasonably, whereas negligence covers cases where one was not aware of (or at least not attending to) a risk but ought to have been.²⁷ Let’s now spell out why the anti-vaxxer’s epistemically unjustified beliefs are (at least in the vast majority of ways of developing the case) either reckless or negligent. If the anti-vaxxer’s beliefs (viz. that the pandemic is a hoax; that the vaccine is designed to implant microchips; etc.) are epistemically unjustified, it follows that there is at least a substantial risk (given the anti-vaxxer’s evidence) that they are false. And if they are false, there is accordingly a substantial risk that by acting in accordance with these beliefs (i.e., advising their relatives not to get the vaccine, etc.), the anti-vaxxer causes disastrous harm. Now, either the anti-vaxxer is aware of this substantial risk that their beliefs are false (and thus, risk occasioning disaster) or she isn’t. If she is aware of the risk of falsity, she is thereby reckless in holding the beliefs despite this. If she isn’t aware of the risk of falsity, then the very fact that the beliefs in question are *unjustified*—depending on some subtle issues²⁸—either

²⁶ Advocates of moral encroachment include Pace (2011), Moss (2018), and Basu (2019), among many others.

²⁷ See e.g. Cowley & Krebs (2020).

²⁸ The subtle issues in question: some think that one can be in a case where one’s evidence is such that one’s belief is unjustified, yet one’s evidence is also such that it doesn’t put one in a position to know, or justifiably believe, that one’s belief is unjustified. (Cf., e.g., Williamson 2014; Lasonen-Aarnio 2020.) If this is possible, then it may be that one can be in a case where there’s a substantial risk (on one’s evidence!) that one’s belief is false, yet it’s not the case that one ought to have been aware of this risk. Whether it is possible is controversial (see e.g. Smithies 2011; Titelbaum 2015 for views that deny that it is), but even those who think it is will allow that these are exceptional cases. It’s also worth noting that the belief may only be negligent if the anti-vaxxer is aware of the harms that might eventuate if the belief is false. We can assume this is very likely in the case at hand.

entails or at the very least strongly suggests that she ought to have been aware of it, in which case she is negligent. And it is morally wrong to be reckless or negligent.²⁹

Second, following our evidence enables us to better achieve what's of value. Again just for illustrative purposes, let's use utilitarianism as a toy theory. Presumably, regardless of whether the utilitarian thinks that you ought to maximize actual utility or expected utility, what they think is ultimately of *value* is actual—not expected—utility. (It would be weirdly fetishistic to think that the *best* state of affairs is one that maximizes expected utility regardless of the well-being of the actual people in it.) But the more epistemically justified our beliefs, the more we can expect our *subjective* (belief-relative) expected utility calculations to converge with what maximizes *actual* utility, and thus to produce this valuable state of affairs. (Again, this does not assume that what we *morally ought* to do in the primary sense is maximize actual utility. The reason to have epistemically justified beliefs can be grounded in fact-relative *value*-facts, not fact-relative *ought*-facts.)

Objection: I am in the course of (eventually) arguing for an evidence-relative view. But mightn't the advocate of the fact-relative view resist my claim that you're morally required to believe what your evidence supports (about morally relevant matters)? Wouldn't applying their view to the doxastic domain naturally lead them to the view that you are, instead, morally required to believe the truth (about morally relevant matters)?³⁰ And won't they be unmoved by considerations of (for example) negligence, since they already think that it's the facts that determine your obligations, not your evidence?

Response 1: first, it's not obvious that the fact-relative theorist will or should be unmoved by considerations of negligence. The fact-relative view says that facts can be eligible to determine your moral obligations whether they are evidentially accessible to you or not. Among these facts are the facts that ground the negligence of any negligent action. So it's unclear why those facts aren't eligible to make an action wrong (because negligent) for a fact-relative theorist.³¹ Now, it's true that a fact-relative *utilitarian* view won't care about negligence, since it only cares about welfare-constituted outcomes. But nonconsequentialist fact-relative views (and perhaps even some nonutilitarian consequentialist ones) may.³² For example, is it really a commitment of the fact-relative view that if you agree to look after my four-year old child, and then leave her alone for an hour, and nothing bad *actually* happens to her, you didn't act wrongly? I don't see why it is, unless objectivism is allied with utilitarianism or some other similar consequentialist view that says that welfare consequences are all that matter.³³

²⁹ Objection: aren't we sometimes blameless for having epistemically unjustified beliefs (cf. Rosen 2003: 63 and n 3, 2004: 301-2)? Reply: yes! On my view, this is just an instance of blameless or non-culpable wrongdoing. One has an *excuse* for the wrongdoing, which precludes one's act—or in this case belief—from being blameworthy, but does not preclude it from being wrong, not from being negligent more specifically. (Indeed, Rosen explicitly notes that negligence need not be culpable: see Rosen 2004: 303).

³⁰ Cf. von Klemperer (ms.). Cf. also Harman (2011: 459), who holds that this is true *with respect to moral truths*.

³¹ Does that mean that the fact-relative view *doesn't* give the verdict I said it did in **Light Switch**? No. The claim is that fact-relative views can count negligence as wrong, not that they'll *only* count cases in which you were negligent as wrong.

³² Thomson (1990) is a case in point. She endorses the fact-relative view (1990: 172-3, 231-2), but nevertheless holds that it's often wrong to impose risks that don't ultimately eventuate in harm (*ibid.*: 243-6). Cf. also Smith (2010: 88) and Fassio (2022).

³³ It's for this reason that I used **Light Switch** and *not* the also-oft-cited miners case (Regan 1980, Parfit 2011, Kolodny & MacFarlane 2010; cf. also Jackson 1991 for similar cases) to illustrate the fact-relative view. It's typically assumed (e.g.

Response 2: the overwhelming consensus view in epistemology is that you're *epistemically* required to believe in accordance with your evidence (not to believe the truth).³⁴ But I think it's quite implausible that we're *epistemically* required to believe what our evidence supports and yet *morally* required to believe the truth. This would set up an utterly pervasive conflict between epistemic and moral requirements on belief. Clause (ii) of Moralistic Evidentialism captures the thought that on the contrary, at least with respect to certain matters our moral obligation is precisely *to be epistemically responsible*.

If Moralistic Evidentialism is true, we have already established that at least some moral principles are evidence-relative, since Moralistic Evidentialism just is an evidence-relative moral principle. But Moralistic Evidentialism is a principle governing belief, which makes it neither a paradigm moral principle generally nor itself the primary locus of debate between fact-relative and evidence-relative views. To get to an evidence-relative view about moral obligations with respect to action, we need to add another premise.

5. Enkrasia

A familiar principle from the rationality literature is

Enkrasia_R. Rationality requires that (if you believe you all-things-considered ought to Φ , you intend to Φ).

Enkrasia_R captures the idea that there is something rationally defective about you if you believe you all-things-considered ought to Φ , but do not intend to Φ . Specifically, Enkrasia_R is typically understood as a requirement of *structural* rationality, or coherence:³⁵ the kind of failure involved in violating it is that of structural *irrationality*, or *incoherence*. Moreover, Enkrasia_R is a “wide-scope” principle, in the sense that the ‘requires’ operator takes scope over the whole conditional, which is stipulatively to be read as material. Wide-scope requirements don’t allow for “factual detachment”: that is, $O(P \rightarrow Q)$ and P don’t entail $O(Q)$. Consequently, Enkrasia_R does not have the bad result that if you believe you ought to do some terrible thing, you’re required to intend to do that terrible thing. All it says is that you’re required to *either* intend to do it *or* not believe you ought to do it. This is compatible with the possibility that further (non-structural) principles determine that, more specifically, you ought to give up the belief rather than form the intention. While Enkrasia_R is not entirely uncontroversial,³⁶ it is still very widely-held.³⁷

by Zimmerman 2006: 337-8; cf. also Prichard 1949/2002: 93) that fact-relative views *must* say that that the right thing to do in the miners case is whatever *actually saves* the miners. But that is only obvious on fact-relative consequentialist views. Other fact-relative views may say that this option is impermissible in virtue of its riskiness.

³⁴ Kelly (2014), for example, writes that this “borders on the platitudinous.” There are claims that *correct* belief is true belief (cf. Wedgwood 2002, 2013; Shah 2003; Gibbard 2005; and many others), but this generally isn’t taken to be the notion at the heart of epistemological theorizing about what you are justified in believing.

³⁵ Cf. n. 7 above.

³⁶ Cf., e.g. Arpaly 2000; Weatherson 2019. Some of the objections that are raised to it can be disarmed by pointing to the features just noted, namely that it is only supposed to be a principle a *structural* rationality, and that it is wide-scope (see Worsnip 2018, 2021).

³⁷ See, among many others, Dancy 1977; Setiya 2007; Brunero 2010; Broome 2013; Coates 2013; Way 2013; Worsnip 2018, 2021.

What’s somewhat less well-explored is the idea that there might also be an enkratic principle that is not a requirement of *rationality* but rather of morality. A moral version of the enkratic principle would differ from Enkrasia_R in at least two ways. First, obviously, it would talk of what morality requires rather than what rationality requires. Second, I think, it should refer to what one believes one *morally* ought to do, rather than what one believes one all-things-considered ought to do. Most theorists agree that it’s not necessarily irrational (or at least, not *structurally* irrational) to fail to intend what you believe you *morally* ought to do. After all, if you think (even wrongly) that non-moral considerations outweigh moral ones, then you may fail to intend what you believe you morally ought to do without any incoherence. Enkrasia_R allows for this by referring to what you believe you all-things-considered ought to do, rather than what you believe you morally ought to do. However, it may nevertheless be *morally* wrong to fail to intend what you believe you morally ought to do.³⁸ The claim that it is, then, would run as follows:

Enkrasia_{M-1}. Morality requires that (if you believe you morally ought to Φ , you intend to Φ).

However, I believe we can make a third revision to the principle once we understand it as a moral (rather than rational) requirement. Specifically, I suggest, the consequent of the embedded conditional can just refer to whether you Φ , rather than whether you intend to Φ . Enkrasia_R refers to your intending to Φ , rather than your Φ -ing, because structural rationality is supposed to be about coherence *between your mental states*. Structural rationality is said to supervene on the mind; if external circumstances interfere with your translating your intention into action, that’s not supposed to be a failure of structural rationality.³⁹ However, almost everyone agrees that *moral* requirements can govern actions, not merely intentions. And indeed, Enkrasia_{M-1} does seem in a respect too weak. If you intend to do what you believe you ought, but never translate this intention into action—if you suffer from what Holton (1999) calls “weakness of will”, in contrast to akrasia—you may be morally criticizable just as much as someone who fails to intend the action at all.⁴⁰ This suggests:

Enkrasia_M. Morality requires that (if you believe you morally ought to Φ , you Φ).

What’s the rationale for Enkrasia_M? To fail to act in line with what one believes one ought to do is to fail to follow one’s conscience. The idea that there is something morally objectionable about failing to follow one’s conscience has a long and venerable history.⁴¹ Failing to follow one’s conscience evinces a kind of morally objectionable hypocrisy. Moreover, if you go against your *beliefs* about what morality

³⁸ Conversely, it *doesn’t* seem like it is always morally wrong to fail to intend what one believes one *all-things-considered* ought to do. If the reasons for which one believes one ought to do the relevant thing have nothing to do with morality, then the failure to intend in line with one’s belief doesn’t seem morally valenced either.

³⁹ See e.g. Broome 2013.

⁴⁰ It may be that if your failure to translate your intention into action is solely due to external impediments, you have an excuse for your failure to act as you believe you ought. But presumably, for *any* moral requirement to act, there will be the possibility of such excuses when external impediments arise. So I don’t think this favors formulating Enkrasia_M in terms of intention rather than action, any more than it does any other moral requirement.

⁴¹ E.g. Ross 1930; Ewing 1947; 1953.

requires (even if these are not in fact true), this manifests a failure of *respect* for morality that is at least *morally* (if, perhaps, not always all-things-considered) deficient.⁴²

Now, there is more than one way to capture the idea that it is morally objectionable to fail to follow one's conscience. A naïve way to capture it would be to say that if you believe you ought to Φ , then you ought to Φ . But this principle has the same bad results that an analogous naïve interpretation of Enkrasia_R does. For example, take **Anti-Vaxxer**. If you completely *unreasonably* believe that you ought to advise your family not to get vaccinated, then the naïve version of the principle says that you ought to advise your family not to get vaccinated. More generally, it gives us the absurd result that it's impossible to hold a false belief about what you ought to do, since any belief about what you ought to do, in effect, makes itself true, just by being believed. This is often called "bootstrapping".⁴³

But to conclude from this that there's nothing morally objectionable about failing to follow one's conscience would be too quick. There are a variety of alternative proposals about how to capture the idea that this is morally objectionable without licensing the implausible results just mentioned. A.C. Ewing (1947; 1953) suggests that we should revise the naïve principle above so that it says, in effect, that if one believes one *objectively* ought to Φ , then one *subjectively* ought to Φ . Yet this works only if the subjective 'ought' both (i) expresses genuine moral obligation *and* (ii) is purely belief-relative—not evidence-relative. And even if both those things are true, it still leaves us with the bad result that merely by believing that you ought to advise your relatives not to get vaccinated, you can make it the case that you really are morally required to do that (even if the moral requirement in question is "subjective"). Conversely, Michael Zimmerman (2006: 336-7) suggests that while what we morally *ought* to do depends on our evidence, we are morally *blameworthy* iff we do what we *believe* is wrong. But this precludes blameworthiness in cases where one doesn't realize that what one is doing is wrong but categorically should have realized this. This seems incorrect.⁴⁴

Enkrasia_M, as (like Enkrasia_R) a wide-scope requirement,⁴⁵ amounts to a different way to capture the moral objectionability of failing to follow one's conscience without any of these problematic results.⁴⁶ It categorically says that you are morally deficient if you fail to do what you believe you ought to do, but because factual detachment fails, it does not allow that merely by believing you ought to do something, you can make it the case that you ought to do it in any sense that involves genuine moral obligation. All that is morally required of you by Enkrasia_M itself is that you *either* form do the thing in question *or* give up your belief that you ought to do it.

⁴² Compare Sylvan (2021) on how failures to respond to *apparent* reasons (where this includes things-believed-to-be-reasons) manifests a lack of respect for reasons.

⁴³ The use of this term with roughly this meaning is (I believe) due to Bratman (1987), who uses it in the context of giving oneself reasons by holding an intention. For its use in the context of a belief about what one ought to do, see e.g. Setiya (2007).

⁴⁴ Zimmerman recognizes this problem (2006: 336 n. 18) and suggests his view will have to be complicated to deal with it. But these complications in turn lose the sense in which it's *always* in a way objectionable to fail to do what one believes one ought to do, even when this belief is itself misguided. The wide-scope formulation can capture that without the bad results of the simpler version of Zimmerman's view.

⁴⁵ Enkrasia_M is proof of concept that even if all requirements of structural rationality are wide-scope, not all wide-scope requirements need be requirements of structural rationality. For other discussions of wide-scope *moral* principles, see Gensler (1985) and Brunero (2024).

⁴⁶ Cf. Dancy (1977), who gives what is to the best of my knowledge the first wide-scope formulation of the enkratic principle and explicitly presents this solution as a rival to Ewing's.

If Enkrasia_M is true, there is an *element* of truth in the *belief*-relative view of moral obligation, since there's at least one moral requirement that makes reference to your *beliefs* rather than the beliefs you ought to have given your evidence, or similar. However, the wide-scope form of the requirement disarms the objections that apply to the belief-relative view proper—which are essentially the same as the objections to the naïve statement of the enkratic principle. And the wide-scope form also means that the principle doesn't conflict with the evidence-relative view: precisely in the cases where your belief is evidentially unjustified, it's the belief you ought to revise in order to satisfy Enkrasia_M.

6. Putting it all together

Moralistic Evidentialism and Enkrasia_M are the main premises in my argument. But I also need two further assumptions.

First, I need the assumption that what we are morally required to believe about what we morally ought to do in particular situations is in some way dependent on what we are morally required to believe about morally-relevant non-moral matters. (Let's call this assumption **Dependence**.) At least at this point, I want to stay neutral on exactly how this dependence works. On one view, what we are required to believe about non-moral matters plus the pure moral *facts* jointly determine what we're required to believe about what we morally ought to do in particular situations. On another view, what we're morally required to believe about non-moral matters *plus* what we're *justified in believing* about the pure moral facts jointly determine what we're required to believe about what we morally ought to do in particular situations.

This difference matters for our resolution of problem #4 for evidential perspectivism, and we'll come back to it. But for the argument at hand, all I need is that what we are morally required to believe about what we morally ought to do in particular situations depends on what we are morally required to believe about morally-relevant non-moral matters—and both views agree on that. Indeed (once we accept that there are moral requirements on beliefs about morally relevant non-moral matters in the first place) I think it would be bizarre to deny this. To deny it would amount to the view that while we are morally required to form beliefs about non-moral but morally-relevant matters that are supported by our evidence, these beliefs should *not* inform our beliefs about what we ought to do in particular situations. That is both very odd in its own right (epistemically and morally justified beliefs about *morally-relevant* matters shouldn't inform our beliefs about what we morally ought to do?), and would also completely undermine the rationale for thinking that we're morally required to follow our evidence about morally-relevant matters in the first place, which is after all that beliefs about these matters *affect our actions*. If some (enkratic) person *didn't* let their beliefs about morally-relevant matters affect their beliefs about what they ought to do in particular situations, then the former *wouldn't* affect their actions.

The second assumption I need (which arguably isn't a premise per se but a logical axiom) is the K axiom of standard deontic logic—also sometimes called “normative detachment”—which says that $O(P \rightarrow Q)$ and $O(P)$ entail $O(Q)$.⁴⁷ [This is distinct from the earlier-discussed and widely-rejected (and not vindicated by standard deontic logic) *factual* detachment principle, which says that $O(P \rightarrow Q)$

⁴⁷ We should read the K axiom as only holding as long as particular “flavor” of ‘ought’ is held constant through the premises and conclusion. That's why it's important that Enkrasia_M is a *moral* principle, not a requirement of rationality.

and P entail $O(Q)$.] The intuitiveness of the K axiom can be brought out by the disjunctive formulation of the wide-scope requirement, $O(Q \vee \sim P)$. If you ought to *either* (make it the case that) Q or (make it the case that) $\sim P$, and but ought (in the very same sense of ‘ought’) to (make it the case that) P, then the option of satisfying the requirement by (making it the case that) $\sim P$ is normatively “out,” so you ought to (make it the case that) Q.

(In fact, I’m not sure I really need that the K axiom holds in *perfect* generality, but just that the inference-pattern it captures is good in the overwhelming majority of cases. For example, it’s really plausible to reason: you ought to [be such that] (if you promise to go to the graduation, you go); you ought to promise to go to the graduation; so, you ought to go.)

Given this, the abstract argument goes as follows. Moralistic Evidentialism states we have a moral obligation to believe what our evidence supports about morally-relevant matters. Dependence says that what we are morally required to believe about what we morally ought to do in particular situations is dependent upon what we are morally required to believe about morally-relevant matters. Given these two assumptions, there will be cases where:

- a) If your evidence supports P, you are morally required to believe that you ought to Φ .
- b) If your evidence supports $\sim P$, you are morally required to believe that you ought not to Φ .

Suppose first that your evidence supports P, and so you are morally required to believe that you ought to Φ . Then, by Enkrasia_M and the K axiom,⁴⁸ you are morally required to Φ . By parallel reasoning, in the case where your evidence supports $\sim P$, you are morally required not to Φ . Hence, whether you are morally required to Φ (or not to Φ) depends on your evidence—even if the *fact of the matter* as to whether P remains constant across both cases. So, what you’re morally required to do (in such a case) depends on your evidence.

It’s hard to fill out this abstract structure with a concrete case without making assumptions about the nature of the dependence relation that Dependence states, which themselves have to be settled at least in part by (further) normative ethical theorizing. But we can give a case that will plausibly do the job on at least most views. Suppose, in an adaptation of **Light Switch**, that you are trying to decide whether to flip a light switch, and that you know that doing so will *either cause* an explosion that kills everyone in the building, *or prevent* an explosion that would kill everyone in the building.⁴⁹ Let P be the proposition that flipping the switch will prevent the explosion, and let Φ -ing be the action of flipping the switch. It’s very plausible, regardless of the details of how will fill out the dependence relation, that given the setup of the case (and Moralistic Evidentialism), (a) and (b) above are both true: if your evidence supports that flipping the switch will prevent the explosion, you are morally required to believe that you ought to flip the switch; but if your evidence supports that flipping the switch will not prevent the explosion (and hence, given the setup, that it will cause the explosion) you are morally required to believe that you ought not flip the switch. Hence, by Enkrasia_M and the K axiom, you ought to flip the switch in the former case but ought not to flip the switch in the latter

⁴⁸ Or, by the inference-pattern that the K-axiom captures, even if this inference pattern goes askew in some funky cases that aren’t this one.

⁴⁹ Let’s also stipulate that you don’t have any weird evidence suggesting that you morally ought to cause the explosion and deaths rather than prevent them.

case—even if we stipulate that the fact of the matter as to whether flipping the switch will prevent the explosion is the same in both cases.

What we have now, if the foregoing argument is sound, is that at least some of our moral obligations depend on our evidence. Now, it's formally consistent with the foregoing argument that we *also* have moral obligations that are determined by the totality of the facts (including those outside our evidence); hence, we haven't established the evidence-relative view in full generality. However, were that the case, we would then face a moral dilemma whenever we had misleading evidence about the facts in question. For our fact-relative obligations in such cases would in effect require us to violate *either* Moralistic Evidentialism *or* Enkrasia_M. If we find it implausible that there are such utterly pervasive dilemmas, we can conclude that we should not allow the additional fact-relative moral obligations that would generate them.

Does that mean that Moralistic Evidentialism and Enkrasia_M constitute the whole of morality? It would be very odd if it did; fortunately, it doesn't. Consider what Moralistic Evidentialism and Enkrasia_M *don't* tell us. They tell us to believe what our evidence supports about morally-relevant matters, and to do what we believe we morally ought to do. What they don't tell us is how to get from our (evidentially-supported) beliefs about morally-relevant (but non-moral) matters to beliefs about what we morally ought to do. All I assumed in my argument was Dependence, which tells us in effect that there should be *some* relationship between the two. We can think of the moral principles that constitute the traditional subject-matter of pure normative ethics as filling in what this relation of dependence is. (If you like, think of it as a function from bodies of evidence or evidentially-supported beliefs (or credences) to claims about what one morally ought to do.) The (evidence-relative) utilitarianism that I've been using as a toy example periodically fills this dependence relation out in one way; other moral theories will do so in other ways.

7. Resolving the problems for evidential perspectivism

I will now explain how the particular version of evidential perspectivism I've defended, and the argument I've given for it, helps us to resolve the problems for evidential perspectivism that I laid out in §3.

First, the lack of consensus on the intuitions about cases. Apart from intuitions, we now have a powerful theoretical rationale for evidential perspectivism. As I noted, the fact-relative theorist might claim that in cases where your evidence is incomplete or misleading, failing to do the (so-called) "objectively right" thing is wrong but excusable. But we now have an argument, beyond appeal to intuition about cases, that in general *performing* this act will be morally wrong. This is because it requires one *either* to violate Moralistic Evidentialism (which is reckless or negligent) or to violate Enkrasia_M (which fails to show due respect for morality).

Second, the sometime unfollowability of the evidence-relative 'ought'. Here I simply grant that the evidentially-constrained 'ought' will not *always* be feasibly followable (for example, in cases where there is insufficient time to work out what your evidence supports, or what you ought to do in line of this), nor will it always be blameworthy to violate it. This isn't a problem for me because I didn't motivate evidential perspectivism by appeal to abstract considerations of followability or "guidance," or the idea that you should *never* be able to act in a way that is wrong but not blameworthy. Indeed,

my view is entirely compatible with the claim that you can act in a way that is wrong but not blameworthy, which I fully endorse.⁵⁰ This may happen in cases where you're misled about what your evidence supports, or when you have limited time to figure it out, and it definitely happens in cases where you lack responsibility for your actions due to lack of control or agency, as in cases of, for example, mental impairment. Rather, I argued for my view from first-order moral principles. In cases where the evidentially-constrained 'ought' is not followable, we may want to advert to a sense of 'ought' *more* subjective than the evidence-relative one, at least for certain purposes (e.g., assigning blame).

Third, the difficulty of accounting for reasons to improve your evidential situation; and relatedly, the (potential) wrongness of acts that result from failing to do so. The view I've defended is very naturally extended to account for reasons to improve your evidential situation. Moralistic Evidentialism tells us that when it comes to morally relevant matters, we have a moral obligation to be epistemically responsible as such. It also says that this involves one's beliefs being properly responsive to one's evidence. However, I think that in certain situations, it also involves gathering more evidence.⁵¹ And the argument for why this is also *morally* required is very similar to the argument for the moral requirement to believe as your evidence supports. Recall that I said (using utilitarianism as an illustration) that in general, we can expect believing as our evidence supports to bring our subjective (belief-relative) expected-utility calculations more in line with what maximizes actual utility. It is also true that gathering more evidence will do the same.⁵² So, at least when (i) we expect the evidence-to-be-gathered to be relatively probative; (ii) it is not too (expectedly) costly to gather and (iii) we have reason to think it bears on something of considerable moral importance, we morally ought to gather more evidence. Thus, the very same moral rationale for believing what our evidence supports extends to improving our evidential situations; and failing to do so will often be morally criticizable, all consistently with the evidence-relative view.

Does this explain the wrongness of *acts* performed on the basis of the evidence one has, when one ought to have gathered evidence that would have favored acting otherwise? Not in itself. However, it may be that we can generalize the argument of this paper to explain this. The idea would be that a similar application of the K axiom can explain why, if you're morally required to gather further evidence, you're then morally required to hold the beliefs that this *enhanced* evidence-set (consisting of your existing evidence plus the evidence you're required to have) supports, and hence (by an application of the original argument) required to act accordingly.⁵³ However, there are tricky issues here about time-indexing: since evidence-gathering can take time, it's not clear whether an obligation to gather further evidence *now* grounds an obligation to act accordingly *now*, rather than at a future time when you've had time to gather the further evidence.⁵⁴ A fuller exploration of these issues will have to wait for another occasion.

⁵⁰ Cf. n. 29 above.

⁵¹ Cf. Worsnip (2019), Flores & Woodard (2023), and McWilliams (2023).

⁵² This is a corollary of Good's theorem (Good 1967). Cf. also Björnsson & Finlay (2010: 16).

⁵³ This would give us something close to the view that Kiesewetter (2011: 16) calls "dynamic perspectivism." See also Ibarra Herrera (ms.) for a somewhat similar, though non-identical, view.

⁵⁴ Thanks to Raúl Ibarra Herrera for discussion of this.

Fourth and finally, the question of what to say about justified but false (pure) *moral* beliefs. As I said when presenting the problem, this is, I think, the most difficult question for evidential perspectivism. With the argument for evidential perspectivism in hand, we're in a position to restate (perhaps with more force) the case for the view that we *are* morally required to act in line with our justified but false moral beliefs. Officially, Moralistic Evidentialism tells us only that we are morally required to believe what our evidence supports about morally relevant non-moral matters. But it is somewhat hard to see why the rationale for this wouldn't also extend to (pure) moral matters. After all, it also seems morally negligent or reckless to form unjustified moral beliefs (given that we act on our moral beliefs). And if we allow that we are morally required to follow our evidence about pure moral matters, then Enkrasia_M and the K axiom yield the result that we're also morally required to act on false moral beliefs that we're epistemically required to have.

Yet, as I said in initially presenting the problem, false moral beliefs are by hypothesis both *false* and *moral* in content. Consequently, it's hard to resist that there is (also) some sense in which you *ought* to act in line with the contrary *true* (yet evidentially-unsupported) principles of pure morality in such cases. To put it more directly: if you in fact morally ought to Φ , then you *morally ought to* Φ —even if your evidence decisively supports the proposition that you morally ought not to Φ .

Of course, it would be convenient (since it would resolve the problem here) if this kind of case were simply not possible—if it were impossible to have evidence that decisively supports false (pure) moral propositions (perhaps because justification for pure moral propositions is a priori and indefeasible).⁵⁵ Unfortunately, I myself don't think that this *is* impossible. So where does this leave us?

My tentative suggestion is that cases where you have this kind of misleading evidence about the true principles of pure morality are moral dilemmas.⁵⁶ You are required to act in line with the true principles of pure morality, but you are also required (by Moralistic Evidentialism and Enkrasia_M) to act in line with the principles of pure morality that your evidence decisively supports. The tragedy of misleading evidence about pure moral matters is that it generates such a dilemma.⁵⁷

Now, you may recall that in §6 I *resisted* the dilemma diagnosis in cases of misleading evidence about non-moral but morally-relevant beliefs. But there is a crucial difference between that case and misleading evidence about pure moral matters. As I stressed in introducing my version of perspectivism, there is no incoherence in thinking that there is *no* sense in which you're *morally* required to respond to non-moral facts outside your epistemic reach. So when it comes to misleading evidence about non-moral facts, there is an easy way to avoid positing a dilemma: we just don't posit any fact-relative moral obligations that would conflict with our evidence-relative ones. By contrast, there *is* incoherence in thinking that there is *no* sense in which you're morally required to respond to *moral* facts (about what's morally required of you!) that are outside your epistemic reach. Thus, if we're required to act on our *justified* false pure moral beliefs, it's much harder to avoid a dilemma diagnosis. In general I think we should avoid positing moral dilemmas when we don't have to, and should avoid

⁵⁵ Cf. Smithies (2023) (and Titelbaum (2015) on analogous claims about rationality. In fact, one argument for this view itself proceeds from an enkratic principle plus an assumption that there are no moral dilemmas, effectively putting my reasoning in this section into reverse.

⁵⁶ Wrinkle: if there are moral dilemmas, the K axiom fails. Yet I relied on it earlier. Here I just have to lean on the point that the reasoning captured by the K axiom is generally good even if it fails to hold in perfect generality.

⁵⁷ Cf. the treatment of misleading higher-order evidence in Worsnip (2018).

positing utterly pervasive moral dilemmas. But it seems to me that in the case of pure moral beliefs, the dilemma diagnosis is the best we have.

With that said, I do think a bit more can be said about such dilemmatic cases to soften any concerns we might have about them. First, such cases don't arise *from the first-person standpoint*.⁵⁸ Precisely because you are misled about the true moral principles, you won't *think* you are in a moral dilemma when this happens: you will simply think the true moral principles *are* as you think they are (given your evidence). So you can't recognize that you're in such a moral dilemma (and consequently, happily, can't feel torn about how to resolve it). Second, and relatedly, I think there is one important asymmetry between your conflicting moral obligations in such a case. You have an *excuse* for failing to comply with the true moral principles about which your evidence has misled you, precisely because your evidence has misled you about them. By contrast, you don't (at least, as the case has been described so far) have an excuse for failing to act in line with the principles of morality that your evidence supports (or, equivalently, for violating Moralistic Evidentialism or Enkrasia_M). (Maybe there's even a sense in which this "resolves" the dilemma, given that one of the conflicting 'ought's can be excusably violated while the other can't be.⁵⁹)

...Unless, of course, your evidence misleads you about Moralistic Evidentialism or Enkrasia_M themselves. Then you do, perhaps, have an excuse for failing to act in line with the principles of morality that your evidence supports. But so it goes. As I have already said, the animating spirit behind my version of evidential perspectivism is not to find moral principles that it's impossible to excusably violate.

One further point is worth making explicit. I have said (as I think we must to avoid incoherence) that when you have misleading evidence about the pure moral principles, you are still required to act in line with the pure moral principles (even though you are *also* required to act in line with the principles that your evidence supports). That means that there are *some* facts—the facts about the pure moral principles—that you are required to take account of even when they are outside your evidence. Does that mean that the fact-relative view is true after all?

No. I framed the fact-relative and evidence-relative views as views about which *non-moral* facts are relevant to determining our moral obligations in particular situations.⁶⁰ And I have not conceded that there are any *non-moral* facts that are outside our evidence but nevertheless relevant to determining our moral obligations. One way to take the upshot of the discussion of the present discussion about (pure) *moral* facts is that if the evidence-relative view is stated in greater generality so that it ranges over the pure moral facts as well, it ends up incoherent or self-defeating. All the more reason not to state it that way. If what I've been arguing in this last section is right, this is not unprincipled. There are real differences between the issues that arise for the non-moral (but morally-relevant) facts and the pure moral facts themselves. We cannot fail to be at least in some sense required to respond to facts that are themselves moral, within our epistemic reach or not. But—if what I've argued in this

⁵⁸ Again, cf. Worsnip (2018). See also Sepielli (2014).

⁵⁹ Cf. Worsnip (2026), where I float a principle called *Excuses Resolve Dilemmas*.

⁶⁰ I've entertained and tentatively endorsed the idea that *Moralistic Evidentialism* is to be extended to pure moral facts. But that is not to say that *evidential perspectivism* is to be extended to pure moral facts.

paper is right—the non-moral facts are not like this. Those outside our epistemic reach cannot play a role in determining our moral obligations.

References

- Arpaly, Nomy (2000). On Acting Rationally Against One's Best Judgment. *Ethics* 110(3), 488-513.
- Basu, Rima (2019). Radical Moral Encroachment: The Moral Stakes of Racist Beliefs. *Philosophical Issues* 29(1), 9-23.
- Björnsson, Gunnar & Finlay, Stephen (2010). Metaethical Contextualism Defended. *Ethics* 121(1), 7-36.
- Bratman, Michael (1987). *Intention, Plans, and Practical Reason*. Harvard University Press.
- Broome, John (2013). *Rationality Through Reasoning*. Wiley-Blackwell.
- Brunero, John (2010). The Scope of Rational Requirements. *Philosophical Quarterly* 60(238), 28-49.
- (2024). Hypocrisy and Conditional Requirements. *Australasian Journal of Philosophy* 102(4), 814-827.
- Clifford, W.K. (1877/1999). The Ethics of Belief. Reprinted in *The Ethics of Belief and Other Essays* (ed. T. Madigan). Prometheus Books.
- Coates, Allen (2013). The Enkratic Requirement. *European Journal of Philosophy* 21(2), 320-333.
- Cowley, Christopher & Krebs, Beatrice (2020). Special Issue on Recklessness and Negligence. *Criminal Law & Philosophy* 14(1), 5-8.
- Dancy, Jonathan (1977). The Logical Conscience. *Analysis* 37(2), 81-84.
- Dowell, J.L. (2013). Flexible Contextualism about Deontic Modals: A Puzzle about Information-Sensitivity. *Inquiry*, 56(2-3), 149-178.
- Ewing, A.C. (1947). *The Definition of Good*. Macmillan.
- (1953). *Ethics*. English Universities Press.
- Fassio, Davide (2022). What the Doctor Should Do: Perspectivist Duties for Objectivists about Ought. *Philosophical Studies* 179(5), 1523-1544.
- Feldman, Richard (1988). Subjective and Objective Justification in Ethics and Epistemology. *The Monist* 71(3), 405-419.
- Flores, Carolina and Woodard, Elise (2023). Epistemic Norms on Evidence-Gathering. *Philosophical Studies* 180(9), 2547-2571.
- Fogal, Daniel & Worsnip, Alex (2021). Which Reasons? Which Rationality? *Ergo* 8(11), 306-343.
- Fox, Philip (2019). Revisiting the Argument from Action Guidance. *Journal of Ethics and Social Philosophy* 15(3), 222-254.
- Gensler, Harry J. (1985). Ethical Consistency Principles. *The Philosophical Quarterly* 35(139), 156-170.
- Gibbard, Allan (2005). Truth and Correct Belief. *Philosophical Issues* 15, 338-350.
- Gibbons, John (2010). Things That Make Things Reasonable. *Philosophy and Phenomenological Research* 81(2), 335-361.
- Good, I. J. (1967). On the Principle of Total Evidence. *British Journal for the Philosophy of Science* 17(4), 319-321.
- Graham, Peter A. (2010). In Defense of Objectivism about Moral Obligation. *Ethics* 121(1), 88-115.

- (2021). Moral Conscientiousness and the Subjectivism/Objectivism Debate about Moral Wrongness. *Oxford Studies in Normative Ethics*, 11, 29-48.
- Harman, Elizabeth (2011). Does Moral Ignorance Exculpate? *Ratio* 24(4), 443-468.
- Hawthorne, John and Srinivasan, Amia (2013). Disagreement Without Transparency: Some Bleak Thoughts. In D. Christensen & J. Lackey, *The Epistemology of Disagreement: New Essays*, 9-30. Oxford University Press.
- Holton, Richard (1999). Intention and Weakness of Will. *Journal of Philosophy* 96(5), 241-262.
- Ibarra Herrera, Raúl (ms.). The Inquiry-Relative View of Moral Rightness and Wrongness.
- Jackson, Frank (1991). Decision-Theoretic Consequentialism and the Nearest and Dearest Objection. *Ethics* 101(3), 461-482.
- Kagan, Shelly (2018). The Paradox of Methods. *Politics, Philosophy, and Economics* 17(2), 148-168.
- Kelly, Thomas (2014). Evidence. *Stanford Encyclopedia of Philosophy*.
- Kiesewetter, Benjamin (2011). "Ought" and the Perspective of the Agent. *Journal of Ethics and Social Philosophy*, 5(3).
- (2017). *The Normativity of Rationality*. Oxford University Press.
- von Klemperer, Caroline (ms.). The Moral Duty to Believe.
- Kolodny, Niko and MacFarlane, John (2010). Ifs and Oughts. *Journal of Philosophy* 107(3), 115-143.
- Lasonen-Aarnio, Maria (2020). Enkrasia or Evidentialism? Learning to Love Mismatch. *Philosophical Studies* 177(3), 597-632.
- (2024). Perspectives and Good Dispositions. *Philosophy and Phenomenological Research* 109(3), 774-798.
- Lord, Errol (2015). Acting for the Right Reasons, Ability, and Obligation. *Oxford Studies in Metaethics* 10, 26-52.
- (2018). *The Importance of Being Rational*. Oxford University Press.
- McGrath, Sarah (2009). The Puzzle of Pure Moral Deference. *Philosophical Perspectives* 23(1), 321-344.
- McWilliams, Emily C. (2023). Evidentialism and Duties to Inquire. *The Philosophical Quarterly* 73(4), 965-982.
- Moore, G.E. (1903). *Principia Ethica*. Cambridge University Press.
- Moss, Sarah (2018). Moral Encroachment. *Proceedings of the Aristotelian Society* 118(2), 177-205.
- Pace, Michael (2010). The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will to Believe'. *Noûs* 45(2), 239-268.
- Parfit, Derek (2011). *On What Matters*, volume 1. Oxford University Press.
- Pittard, John (2023). The Many 'Oughts' of Deliberation. *Philosophical Studies* 180(9), 2617-2637.
- Pittard, John and Worsnip, Alex (2017). Metanormative Contextualism and Normative Uncertainty. *Mind* 126(501), 155-193.
- Prichard, H.A. (1949/2002). Duty and Ignorance of Fact. Reprinted in his *Moral Writings* (ed. J. McAdam), 84-101. Oxford University Press.
- Rosen, Gideon (2003). Culpability and Ignorance. *Proceedings of the Aristotelian Society* 103, 61-84.
- (2004). Skepticism about Moral Responsibility. *Philosophical Perspectives* 18, 295-313.
- Ross, W.D. (1930). *The Right and the Good*. Oxford University Press.
- (1939). *Foundations of Ethics*. Oxford University Press.

- Scanlon, T.M. (2008). *Moral Dimensions: Permissibility, Meaning, and Blame*. Harvard University Press.
- Sepielli, Andrew (2014). What to Do When You Don't Know What to Do When You Don't Know What to Do... *Noûs*, 48(3), 521-544.
- Setiya, Kieran (2007). Cognitivism about Instrumental Reason. *Ethics* 117(4), 649-673.
- Shah, Nishi (2003). How Truth Governs Belief. *The Philosophical Review* 112(4), 447-482.
- Smith, Holly M. (2010). Subjective Rightness. *Social Philosophy and Policy* 27(2), 64-110.
- (2014). The Subjective Moral Duty to Inform Oneself before Acting. *Ethics* 125(1), 11-38.
- (2018). *Making Morality Work*. Oxford University Press.
- Smithies, Declan (2011). Moore's Paradox and the Accessibility of Justification. *Philosophy and Phenomenological Research* 85(2), 273-300.
- (2023). The Problem of Morally Repugnant Beliefs. *Oxford Studies in Metaethics* 18, 218-241.
- Srinivasan, Amia (2015). Normativity Without Cartesian Privilege. *Philosophical Issues* 25, 273-299.
- Sylvan, Kurt L. (2021). Respect and the Reality of Apparent Reasons. *Philosophical Studies* 178(10), 3129-3156.
- Thomson, Judith Jarvis (1986). Imposing Risks. In her *Rights, Restitution, and Risk* (ed. W. Parent), 173-191. Harvard University Press.
- (1990). *The Realm of Rights*. Harvard University Press.
- Titelbaum, Michael G. (2015). Rationality's Fixed Point (or: In Defense of Right Reason). *Oxford Studies in Epistemology* 5, 253-292.
- Way, Jonathan (2013). Intentions, Akrasia, and Mere Permissibility. *Organon F* 20(4), 588-611.
- Weatherson, Brian (2019). *Normative Externalism*. Oxford University Press.
- Wedgwood, Ralph (2002). The Aim of Belief. *Philosophical Perspectives* 16, 267-297.
- (2007). *The Nature of Normativity*. Oxford University Press.
- (2013). Doxastic Correctness. *Proceedings of the Aristotelian Society Supplementary Volume* 87, 217-234.
- Williams, Bernard (1981). *Ought* and Moral Obligation. In his *Moral Luck*, 114-123. Cambridge University Press.
- Williamson, Timothy (2014). Very Improbable Knowing. *Erkenntnis* 79(5), 971-999.
- Worsnip, Alex (2018). The Conflict of Evidence and Coherence. *Philosophy and Phenomenological Research* 96(1), 3-44.
- (2019). The Obligation to Diversify One's Sources: Against Epistemic Partisanship in the Consumption of News Media. In J. Saunders & C. Fox (eds.), *Media Ethics, Free Speech, and the Requirements of Democracy*, 240-264. Routledge.
- (2021). *Fitting Things Together: Coherence and the Demands of Structural Rationality*. Oxford University Press.
- (2026). Evidence-Coherence Conflicts Revisited. In N. Hughes (ed.), *Epistemic Dilemmas*. Oxford University Press.
- Zimmerman, Michael J. (2006). Is Moral Obligation Objective or Subjective? *Utilitas* 18(4), 329-361.
- (2008) *Living with Uncertainty*. Cambridge University Press.
- (2014). *Ignorance and Moral Obligation*. Oxford University Press.