

Epistemic Oughts and Epistemic Dilemmas

The following chapters examine the theoretical upshots of the positive epistemological view proposed in this book. The account developed so far delivers the result that epistemic justifiers constitute epistemic oughts. In this chapter, I discuss the worry that such accounts threaten to give rise to widely spread epistemic dilemmas between paradigmatic epistemic norms. I argue for a modest scepticism about epistemic dilemmas. In order to do that, I first point out that not all normative conflicts constitute dilemmas: more needs to be the case. Second, I look into the moral dilemmas literature and identify a set of conditions that need to be at work for a mere normative conflict to be a genuine normative dilemma. Last, I argue that while our epistemic life is peppered with epistemic normative conflict, epistemic dilemmas are much harder to find than we thought.

10.1 Obligations to Believe and Epistemic Dilemmas

My account takes epistemic justification to be epistemic obligation. Nevertheless, it does not follow from my view that we are obliged to believe all of the things that, for example, happen in our visual fields at the same time. That is because what grounds obligations to believe is being in a position to know, and being in a position to know is limited quantitatively by our cognitive capacities.

We have also seen that, on my account, paradigmatic epistemic norms such as K_p will not come into conflict with zetetic norms, or norms of inquiry and gathering evidence. The question that arises at this point is: how about paradigmatic epistemic norms themselves? Will they not generate normative conflicts and, indeed, normative dilemmas?

First, recall that the account predicts that normative priority will be decided by availability rankings. As such, short of too many facts (i.e. more than I can uptake at once) enjoying the precise same availability ranking,

the view will not result in normative dilemmas. Second, although epistemic dilemmas have been hotly discussed in recent epistemology, I remain unconvinced of their in-principle possibility. The following sections, indeed, will argue for a (modest) scepticism about epistemic dilemmas. If I am right, while our epistemic lives – in line with our normative lives more generally – encounter decidable epistemic conflict often, epistemic dilemmas are rather hard to come by.

10.2 What Dilemmas Are Not

What is a normative dilemma? On a first approximation, it seems plausible that one is facing a dilemma just in case two courses of action and two courses only are available to one, and whichever one chooses, one finds oneself in normative breach: there's no good way out, as it were. This seems promising. Let's spell it out:

Normative dilemma #1 (ND1): A state of affairs such that a subject S has only two available courses of action, both of which imply norm violation.

ND1 is false – and widely acknowledged¹ to be false – in that it is too inclusive: it defines conflicts rather than dilemmas. Not all normative conflicts are dilemmas: there are several well-theorised phenomena that prevent garden variety normative conflicts from becoming full-blown dilemmas.

First, in cases in which one norm overrides the other, conflict is present, whereas there's no dilemma as to the best action to pursue: it is the one recommended by the overriding norm. Indeed, it lies in the meaning of 'overriding' that such cases are cases of normative conflict without a dilemma. That's what it is for a norm to override another: to come into conflict with it and take precedence. If it takes precedence, there's no dilemma left to face the subject of the normative constraints in question: again, I face a moral conflict but not a moral dilemma when I decide to save the drowning child even at the cost of breaking my promise to meet you for lunch at noon. The norm of saving lives conflicts with the one of promise-keeping and renders one course of action, and only one, permissible – no dilemma here. Our definition needs tightening up if we are to distinguish dilemmas proper from mere normative conflicts. In addition to the conditions stipulated by ND1, at a minimum, it also needs to be the case that neither of the normative constraints are overridden.

¹ For an excellent overview of the relevant literature, see McConnell (2018).

Second, normative requirements can be overridden, but they can also be undercut. The difference between undercutting and overriding is that, roughly, while in cases of overriding we get reasons for and against a course of action ϕ , and the reasons against ϕ -ing are, for example, weightier than the reasons for ϕ -ing – such that the latter get overridden by the former – in cases of undercutting the counter-reasons speak, in the first instance, against the normative strength of the reasons in favour of ϕ -ing rather than directly against ϕ -ing. Your testimony that the train leaves at 8 a.m. is reason for me to go to the station before 8 a.m. My finding out that you are a compulsive liar undercuts this reason and renders it normatively inert. There is normative conflict between the two reasons, for sure. But the conflict fails to result in a dilemma: I know precisely what to do in this situation (i.e. not base my action on trust in your testimony).

Our definition needs to be tightened up if it is to distinguish between mere normative conflicts and normative dilemmas proper. At a minimum, on top of the conditions stipulated by ND₁, we need to add an anti-overriding and an anti-undercutting condition. Here is a second pass:

Normative dilemma #2 (ND₂): A state of affairs such that a subject S has only two available courses of action, both of which imply *active* norm violation.

In ND₂, what is meant by ‘active norm’ is a norm that remains efficacious within the context after all things normative are considered. What ND₂ implies is that, for a dilemma to be instantiated, (1) the two normative constraints need to be equally weighty (on pain of overriding taking place) and (2) it should not be the case that the one sheds doubt on the normative credentials of the other.

On closer inspection, though, ND₂ is still too broad: sometimes, two active (equally weighty, non-undercut) norms can come into conflict, while a dilemma is not instantiated in virtue of the fact that one takes qualitative precedence over the other. Of course, overriding and undercutting are also ways in which norms can take precedence; they are not the only ways, however. Overriding is a quantitative matter: weightier normative constraints prevail. Precedence relations, however, can also be qualitative: they can be exhibited between active, non-overridden norms as well. A paradigmatic such case is one in which one of the norms is derivative of the other. Here is such a case:

Promise-breaker George: George is a promise-breaker: whenever possible, he will reliably break his promises to others. One day,

George promises his colleague, Anna, to call her on Thursday at precisely 12 o'clock, and, as per usual, he doesn't care much about keeping his promise. On Thursday, he looks at his watch, comes to believe it's 11.45 a.m., and decides to take a nap before making the call, thinking that he will be at most thirty to forty minutes late. Luckily, though, George's watch is broken: it's actually only 10.45 a.m. After taking his nap, George ends up calling Anna at precisely 12 o'clock.

Promise-breaker George is in breach of a bunch of conceivable norms in this scenario (Williamson forthcoming): he has a bad disposition – he is a promise-breaker; he acts in ways that would, had his watch not been broken, have resulted in breaking his promise; in general, he seems to be rather inconsiderate and untrustworthy.² But here is one thing that cannot be said about George: that he broke his promise to Anna. Morally lucky George did no such thing. Indeed, he called Anna at precisely 12 o'clock as promised.

What is happening in this case is an instance of breach of a number of norms that derive from what we may call, following Williamson (forthcoming), the 'primary' norm of promise-keeping ('keep your promises!'; e.g. 'don't be a promise-breaker!', 'don't act like a promise-breaker would!') without breach of the primary norm.

What is crucial to note about this case is that George cannot comply with both the primary norm of promise-keeping – 'keep your promises' – and, for example, the derived norm 'do what a promise-keeper would do!'. Meeting one will ensure he is in breach of the other. In that, one would think, the case looks initially promising for a dilemmatic case. Had George not been in breach of some of the derived norms, he would have been the victim of moral bad luck and would have ended up in breach of the primary norm at stake: had he, for example, called Anna at what he thought – according to his broken watch – was 12 o'clock, he would have failed to keep his promise (since he would have, in fact, called her at 11 o'clock). Last but not least, note also that this is not a case of either overriding or undercutting: the norms at stake here – both primary and derived – are standing, active norms. It's easy to see this from the fact that, whichever of them George would be in breach of, he would intuitively be the proper subject of blame. It's worthy of blame to break your promises, but so is to be the kind of person who would do so and to act accordingly.

² See Kelp and Simion (2023a) for an account of trustworthiness in terms of a disposition to comply with one's obligations.

Still, George is not faced with a moral dilemma here: the primary norm of promise-keeping, as its description suggests, has priority. To see this, note that we describe the case as one of moral *good* luck. Had the norm of promise-keeping not taken precedence, it's a mystery why we wouldn't describe the case in neutral terms (should the norm of promise-keeping be as stringent as the norm requiring one to be the kind of person who keeps promises, etc.) or negative terms (should what we have dubbed 'derived' norms take precedence).

What George's case suggests is that we need a further (and final) restriction on our account of dilemmas: a state of affairs will qualify as a dilemma if and only if a subject S has only two available courses of action, both of which imply *active* norm violation, where neither of the norms at stake is derivative of the other. Here is a simpler way to put this:

Genuine normative dilemma (GND): A state of affairs such that a subject S has only two available courses of action, both of which imply norm violation, where neither of the norms at stake takes precedence over the other.

In turn, as we have seen, taking precedence can come in many shapes, including overriding, undercutting, or normative primacy. William Styron's (1979) *Sophie's Choice* presents a useful example of such a genuine moral dilemma. Sophie Zawistowska has been asked to choose which of her two children, Eva or Jan, will be sent to the gas chamber in Auschwitz. An SS doctor, Fritz Jemand von Niemand, will grant a dispensation to only one of Sophie's children. If she does not choose which one should live, Dr von Niemand will send both to their deaths.

In Sophie's case, the moral norm asks of her to make a choice rather than not; otherwise, the worst scenario obtains: both children will die. So withholding from acting is not an option. The two options are: sacrifice Eva, or sacrifice Jan. Both options represent breaches of standing, non-overridden (the two options are equally bad, no norm is weightier than the other), non-undercut moral (and prudential) norms, neither of which takes priority over the other. Indeed, plausibly, the moral norm at stake on both horns of the dilemma is one and the same: 'don't put the life of your child in danger!'. In this, *Sophie's Choice* is a GND.

10.3 Epistemic Non-dilemmas

This section puts the results above to use: it looks at several cases featured in the epistemological literature as alleged examples of epistemic dilemmas

and argues that they are garden variety normative conflicts rather than genuine dilemmas.

Let's start with a very straightforward case:

Rebutting defeat: My four-year-old hears me coughing and tells me I have a cold. My doctor disagrees: it's bronchitis.

Here's a garden variety epistemic normative conflict that is not a normative dilemma. Rather, it's a straightforward case of normative overriding: I have stronger reason to trust my doctor's testimony than my four-year-old son's. Indeed, note that epistemological terminology implies that this is a non-dilemmatic normative conflict: after all, this is a classic case of full rebutting defeat acting against the epistemic reason provided by my son's testimony. The presence of full rebutting defeat, however, implies that the reasons against my belief that it's a cold are weightier than those in favour: otherwise, full defeat would not be instantiated. The presence of defeat, then, precludes this variety of normative epistemic conflict from being a genuine dilemma. Indeed, this is, of course, a mere epistemic incarnation of the case in which I am late for lunch because I stop to save the child: a classic case of normative overriding.

What about a case where two equally reliable, equally trustworthy, etc., sources (plug in your favourite view of the epistemology of testimony) offer conflicting testimonies? No problem at all: uncontroversially, the epistemically correct thing to do is to suspend/withhold belief³; again, there is no epistemic dilemma here. Indeed, strictly speaking, we should expect epistemology to be, if anything, most often the proper home of normative trilemmas rather than dilemmas: after all, in epistemology, there's always the possibility to suspend belief.

So far so good: one wouldn't expect much in the way of controversy to be triggered by this fairly straightforward diagnosis of rebutting defeat.

But here is a type of case hotly discussed under the heading of an epistemic dilemma, starting back in the early 1990s:

Evidence-undermining belief: As S considers some proposition, *p*, it is clear to S that an effect of S's believing *p* would be to undermine the evidence S has, which otherwise is sufficient epistemic reason for S to believe *p*.

³ The difference between withholding and suspending belief will be of no consequence throughout this chapter. I will therefore use them interchangeably.

Evidence-undermining belief is a type of case extensively discussed in several epistemological works (e.g. Conee 1987, 1992, Sorensen 1987, Richter 1990, Foley 1991, Kroon 1993, Odegard 1993). Here, for example, is Odegard's take on the normative landscape present in this type of case:

Clearly we should not deny the belief, since this would be to deny a belief for which we have adequate evidence, both prior to adopting a position on it and when we adopt a position on it. But we should not affirm the belief either, for this would be to affirm a belief for which we would not have adequate evidence when we held it [. . .]. Yet it can seem that we should not withhold on the belief either, since in withholding on it we fail to adopt a belief for which we have adequate evidence when we consider it for adoption. So it can seem that whatever happens, we do something that we should not do. (Odegard 1993, 161)

The discussion in the Section 10.2, however, should by now have made it clear that Odegard's trilemmic diagnosis here is mistaken. What we have here is a straightforward case of normative undercutting: the normative force of the evidence that the subject has for thinking that p is undercut by the evidence that the subject has for thinking that as soon as they adopt a belief in p , this would undermine the normative strength of the evidence for p . This is a straightforward case of undercutting defeat. Indeed, here is Conee's diagnosis along the same lines:

When believing would result in a loss of crucial evidence for the believed proposition, adopting the belief would not bring about knowledge of the proposition. Foreseeing this sort of loss excludes having an epistemic reason to believe when contemplating the proposition. (Conee 1993, 478)

If this is a case of undercutting defeat, however, it cannot, in principle, constitute a normative dilemma, for roughly the same reason why cases of rebutting defeat cannot constitute dilemmas: that is what it means for evidence to be defeated – it is for it to lose its initial normative strength. Fully undercut evidence no longer supports one's belief in the target proposition. The case cannot be at the same time one of full undercutting defeat and a dilemma.

How about partial undercutting defeat? Can't it be that the dilemma arises when the undercutter only partially affects the first-order evidence? Consider:

Logic problem: Anna is a logic student who is evaluating a tautology (T). Anna is certain that (T) is true. However, her logic professor, Chad, then tells her that before she began the exam, she was slipped a reason-distorting

drug that impairs one's ability to solve logic problems; those who are affected by the drug only reach the right conclusions 50 per cent of the time. As it turns out, though, unbeknownst to both Anna and Chad, the drug was just a placebo and Anna's logic reasoning abilities were not affected in the least. (Adapted from Leonard 2020)

Chad's testimony provides Anna with higher-order evidence to the effect that there is a 50 per cent chance that she botched her assessment of what this first-order evidence actually supports. Her first-order evidence is thus partially undercut. What is Anna supposed to believe? According to some,⁴ Anna is faced by a dilemma that requires a lot of epistemic fine-tuning to explain away.

However, note that, again, insofar as we accept the case as one of undercutting defeat, it can't be that this is a dilemma rather than a mere normative conflict (and if we don't thus accept it, no dilemma arises): its being a case of partial defeat implies that the higher-order evidence partially neutralises the normative force of the first-order evidence. The fact that the defeat is merely partial does nothing to change this: by the way the case is built, at least on a first approximation, Anna is left with 50 per cent first-order normative support for her belief that (T). If so, there is no dilemma here: Anna should suspend belief, since she has equal support for (T) and non-(T).

To see this further, let's see what would have to be the case for this to be a genuine dilemmatic case. For a genuine dilemmatic normative conflict to be instantiated, we would have to think that something like the following principle (Worsnip 2018) holds:

Possibility of Iterative Failure (PIF). It is possible that:

- i. S's evidence supports $D(p)$; and
- ii. S's evidence supports believing that her evidence does not support $D(p)$,

where $D(p)$ is a possible doxastic attitude for a subject S towards a proposition p . I find PIF implausible, for the following reason (which I expanded on in Chapter 8): evidence about what evidence supports will normatively affect what evidence supports, in one way or another: sometimes, when weightier than the first-order evidence, it will defeat its normative strength. Andy's testimony that p : 'the train leaves at 8 a.m.' is a reason for me to believe that the train leaves at 8 a.m. Testimony from the much more reliable (trustworthy, etc.) Mary that Andy is a compulsive

⁴ See Leonard (2018) for an overview and discussion.

liar is a reason for me to believe that Andy's testimony is less weighty than I thought and thus lower my confidence in p . This need not always be the case, of course. It may be that defeat goes the other way – in cases in which the first-order evidence is weightier. My a priori justification that there are no round cubes will likely defeat my four-year-old's testimony that I'm confused, since he just saw a round cube at his friend's house. At other times, the two sources can also be equally weighty, in which case, again, the proper thing to do is to suspend. The important point, though, is that higher-order evidence interacts normatively with first-order evidence, which renders PIF implausible.⁵

Do we have any reason to believe PIF to be true, in spite of this prima facie plausible evidentiary situation? Alex Worsnip argues that we do; according to Worsnip, rejecting PIF commits one to an implausibly strong claim about justification: denying (PIF), according to him, requires denying that one can have, all things considered, misleading evidence about what one's evidence supports; in other words, justified false beliefs about what one's evidence supports are impossible. That is a strong claim, as any claim that a particular kind of justified false belief is impossible would be (Worsnip 2018).

Let's state this clearly. According to Worsnip, the following claim holds:

Non-PIF implies no justified false beliefs about evidential support (NJFBES): If non-PIF, then one cannot have justified false beliefs about what one's evidence supports.

I agree with Worsnip that NJFBES is implausibly strong.⁶ I disagree, though, with the claim that the denial of PIF implies it. In particular, denying PIF is perfectly compatible with having a justified false belief that your evidence *does* support p . Rather, what the denial of PIF implies is the weaker:

No justified false beliefs about lack of evidential support (NJFBLES): One cannot have a justified false belief that one's evidence does not support p .

NJFBLES might be hard to recognise at first, but, as opposed to its more ambitious cousin NJFBES, it is not a particularly controversial claim: it's

⁵ See Chapter 8 for an extensive critical discussion of scepticism targeting this claim.

⁶ Note, though, that traditionalist evidentialists might have to accept it. If one's justification is strictly a function of one's evidence, then it seems to follow that one cannot have justified false beliefs about what one's evidence supports.

merely stating that undercutting defeat via higher-order evidence is a genuine phenomenon. As soon as you have justification for thinking your evidence does not support p , either the higher-order justification negatively affects the normative strength of your first-order evidence, making it true that you don't have evidential support for p , or, if weightier, your first-order evidence affects the normative strength of the higher-order justification, making it false that you are justified. If you think there is such a thing as undercutting defeat, you hold NJBLES to be true (and for people who don't, see the discussion in Chapter 8 on scepticism about the defeating power of higher-order evidence, as well as the discussion below on knowledge-first, level-splitting views). If so, the fact that denying PIF does imply NJFBLES (but not NJFBES) is not a problem for denying PIF but at worst a natural feature, and it is more likely a theoretical virtue of prior plausibility (given the widely spread popularity of undercutting defeat).

To sum up: for all of the above cases, I have argued that if we accept that they are cases of defeat, their being epistemic dilemmas becomes an in-principle impossibility, since their being cases of defeat implies that they are cases of either normative overriding or normative undercutting. How about if one wants to deny the very idea of defeat? Here is another type of case featuring what seems to be undercutting defeat that has been discussed in more recent literature:

Maths: A competent mathematician has just proved a surprising new theorem. She shows her proof to several distinguished senior colleagues, who all tell her that it involves a subtle fallacy. She cannot quite follow their explanations of her mistake. In fact, the only mistake is in their objections, obscured by sophisticated bluster; her proof is perfectly valid. (Williamson forthcomingb, 1)

In this case, too, if we accept that what is going on is undercutting defeat, we are left with mere normative conflict without normative dilemma: depending on the weight of the colleagues' testimony, the mathematician's first-order support for the theorem will be more or less diminished. The amount of warrant left will support either belief (if only marginally affected) or disbelief (if seriously affected), or else suspension.

Several authors in the knowledge-first camp, though, argue against undercutting defeat for knowledge. According to people like Maria Lasonen-Aarnio (2014) and Tim Williamson (forthcomingb), insofar as our mathematician knows that the theorem in question holds, misleading higher-order evidence will have no impact on the normative credentials of their belief: they should hold steadfast. In turn, these philosophers explain the intuition of the impermissibility of such dogmatic doxastic behaviour

via appeal to epistemic blameworthiness. According to Lasonen-Aarnio, the intuition that dogmatism is suspicious doxastic behaviour even in the presence of knowledge is to be explained by the fact that ignoring evidence is, generally speaking, a bad epistemic disposition that is worthy of blame. As such, while our mathematician is not in breach of the norm of belief in this case – since they are a knower – they are blameworthy for displaying a bad epistemic disposition in ignoring available evidence.

Does this take on these cases create problems for our diagnosis of them as non-dilemmatic normative conflicts? The answer is ‘no’. Insofar as one holds that there is some sort of priority ordering between the two norms coming into conflict – the knowledge norm of belief on the one hand and the norm prescribing against a disposition to ignore evidence on the other – the case is not an epistemic dilemma. Recall the case of George the promise-breaker: just like in that case, insofar as one takes one of these norms to have primacy over the other, GND is not instantiated.

Now, it is plain to see that according to the no-defeat champions, the knowledge norm takes primacy in Maths: first, because they hold that the mathematician should hold steadfast in this case – which suggests that the knowledge norm takes primacy over the dispositionalist norm – and second because they hold that the mathematician is in mere blameworthy norm compliance rather than in genuine norm violation. If so, by the lights of this variety of undercutting defeat deniers, there will be no dilemma instantiated in this case.

We have seen that both champions and foes of undercutting defeat will have to deny that cases like Maths instantiate epistemic dilemmas. Note, though, that there is still a bit of theoretical distance between this result and an in-principle impossibility of cases like these being dilemmatic: after all, there is one possibility left in the logical space. One could deny undercutting defeat and at the same time hold that the first- and second-order norms in this case have equal normative strength: none takes primacy. If so, one would think, we would have an instance of GND in this case: an epistemic *Sophie's Choice*.

Fortunately, though, that's not quite right. Our epistemic lives are easier than our moral lives: what is often an epistemically available option, but is not always a morally or prudentially available option, is suspending. By stipulation, Sophie does not have the option not to make any choice between her children: if she refuses to choose, they will both be killed. In cases of epistemic conflict, however, suspending belief is often an available option. As such, mere normative strength parity will not be enough to generate a dilemma.

This will be the case in both cases of (alleged) undercutting and rebutting defeat: for the theorist who rejects defeat and upholds normative parity, what is going on in these cases is a conflict between two equally weighty norms – one requiring belief and one disbelief in the relevant target proposition. If so, epistemology has an easy answer to these cases: the subject must, *ceteris absentibus*, suspend.

What the discussion so far suggests is that epistemic dilemmas are hard to come by. What we would need to generate an epistemic *Sophie's Choice* is, for example, an equally weighty reason against believing that p and believing that non- p and an even stronger reason against suspending belief. Here it is:

Genuine epistemic dilemma (GED): A state of affairs such that believing that p , believing that non- p , and suspending on p all imply epistemic norm violation, where the norm forbidding one of the three options is weightier than the remaining two and neither of the remaining two norms takes precedence over the other.

Alternatively, we can also have a genuine epistemic trilemma, should the norms in question be equally weighty:

Genuine epistemic trilemma (GET): A state of affairs such that believing that p , believing that non- p , and suspending on p all imply epistemic norm violation, where none of the norms at stake take precedence over the others.

Note how far we've come from our first pass at isolating dilemmas proper from mere normative conflicts: normative conflicts are ubiquitous in epistemology; di/trilemmatic conflicts, however, are less so, if they exist at all.

Could we get something like GED/GET in our epistemic life? One thing to notice, from the start, is that what we would need is a proper *epistemic* reason against suspending. Stipulating that you're bound to either believe or disbelieve because a villain is holding a gun to your head and threatens to kill you if you suspend, even though your evidence equally supports p and non- p , will not generate an epistemic dilemma, but rather an inter-normative conflict with an easy way out: all things considered, you should randomly believe whatever just to save your life. Epistemically, though, you should suspend.

I would like to end this chapter on a more optimistic note than I have proceeded with so far: I would like, that is, to propose two cases that, at least at first glance, look to me like better candidates for an epistemic

di/trilemma than what we have been looking at so far. I am not myself convinced that they will hold water ultimately (which is why I dub them, modestly, ‘attempted epistemic di/trilemmas’). But it does seem to me, in the light of the results in this chapter, that they stand a better chance at instantiating di/trilemmatic normative conflict proper than the cases that we have been looking at. Because of this, I think they are worth putting on the table for further discussion. Here they are:

Attempted epistemic trilemma (AET): Mary, John, and Anna are equally reliable, equally trustworthy testifiers, and you know them to be such (again, plug in whatever else you need to instantiate epistemic justification on your favourite view of testimony). Mary tells you that p : the train leaves at 8 a.m. John tells you that non- p : the train does not leave at 8 a.m. Anna tells you that you don’t have equally weighty evidence for p and non- p (alternatively, Anna tells you that it’s epistemically impermissible for you to suspend belief on whether the train leaves at 8 a.m.).

And, correspondingly:

Attempted epistemic dilemma (AED): Mary and John are equally reliable, equally trustworthy testifiers, and you know them to be such. Mary tells you that p : the train leaves at 8 a.m. John tells you that non- p : the train does not leave at 8 a.m. Anna is the most reliable (trustworthy, etc.) testifier you know. Anna tells you that you don’t have equally weighty evidence for p and non- p (alternatively, Anna tells you that it’s epistemically impermissible for you to suspend belief on whether the train leaves at 8 a.m.).

There are a few things to notice about these cases. First, note that the cases need not be spelled out as featuring testimony; the choice here is driven by convenience. Parallel cases can be described with any other sources of knowledge. Nor does it have to be the case that one and only one type of source is at stake: a combination would do, too. Second, about AED: it is meant to be the epistemic equivalent of *Sophie’s Choice* structurally. Third, note that AED and AET are only di/trilemmas if we assume that an undercutting defeat-denying view is false, and thus that the higher-order evidence provided by Anna affects the justification you get from the first-order evidence generated by Mary and John. Otherwise, the case will be one of permissible suspension, and thus no dilemma will be instantiated.

Are AED and AET genuine epistemic di/trilemmas? Again, I’m not fully convinced: it may depend on what the correct view of evidential

weight will be (e.g. the correct view of evidential weight might make it such that what one should do in these cases is suspend on everything: p , non- p , and the issue of what your evidence supports). I do believe, though, that these cases are worthy of serious attention, in that, as opposed to other cases that are historically popular in the literature, they do instantiate a di/trilemmatic structure proper: it looks as though, that is, whatever one decides to do – doxastically speaking – in these cases, one is in breach of equally strong, standing norms, neither of which takes priority over the other. (Compatibly, of course, structure might not be all that there is to epistemic dilemmatic conflict.)

10.4 Conclusion

I have defended modest scepticism about epistemic dilemmas: they're hard to find. My scepticism, to be clear, only falls short of being radical insofar as the attempts I made at mimicking a *Sophie's Choice* structure for the epistemic – or similar attempts – can be made to work. I am not myself convinced that they will, however. If they turn out to fail, I want to claim that we have reason to be very pessimistic about the in-principle possibility of an epistemic dilemma. I have also argued that an account like mine, taking justifiers to be obligations, will not deliver the theoretically suspicious result that the epistemic domain is peppered with dilemmas, but rather the more modest result that, just as in other normative domains, the epistemic sometimes faces us with normative conflict: sometimes, I can't make it to lunch in time and save the child from drowning – something has to give. Epistemically, things look pretty similar.