

n-1 Guilty Men¹

Julien Dutant and Clayton Littlejohn

Abstract: We argue that there is nothing that can do the work that normative reasons are expected to do. A currently popular view is that in any given situation, a set of normative reasons (understood as a set of facts, typically about the agent's situation) always determines the ways we prospectively should or should not respond. We discuss an example that we think shows no such collection of facts could have this normative significance. A radical response might be to dispense with reasons and explain the facts about the ways we should respond using something else entirely (e.g., expected value). We think a more fruitful approach is to explore the possibility of using decision-theoretic tools in a reasons-centred framework.

Keywords: blame, epistemic filter, expected value, knowledge, normative reason, objective 'ought', prospective 'ought', rationality, uncertainty.

0. The Difficult Weekend

Agnes looks after n kids (i.e., 1, 2, 3, ..., Inge, ..., $n-3$, $n-2$, $n-1$, n). During one difficult weekend, she came to suspect something each of them has done something bad. She suspected that $n-9$ stole something from 7, that 7 bullied $n-14$, and so on. When she suspects that a kid breaks a rule, she gives them a fair hearing. When she reaches a verdict, she only judges that a kid is guilty using the right (?) kind of evidence.² Being the very model of epistemic rationality and moral conscientiousness, she only believes a kid to be guilty when the evidence is sufficient, only blames if found guilty

¹ Title inspired by Volokh (1997). For Volokh, 'n' is the number we solve for if we vindicate or improve upon Blackstone's suggestion that it would be better to let ten guilty men go free than convict one innocent person. (According to Blackstone, $n = 10$.) For us, 'n' is the number of people being punished. Our tale involves kids, but kids can be any age as some of our parents will call us 'kids' well into our 40s. We needed a case involving blame and sanction in a non-institutional setting. Some readers might prefer to think of the kids as adults.

² There has been considerable discussion recently of what the right kind of evidence might be. We make few assumptions about this apart from the assumptions that (a) we can have sufficient evidence to believe that someone is guilty of and responsible for some wrongdoing and (b) that naked statistical evidence is not the right kind of evidence to rationalise blame or punishment. These assumptions are not completely uncontroversial, but they are not terribly controversial. These assumptions are often taken for granted in discussions of the standard of proof.

of breaking a rule that should be followed, and only dishes out fitting punishments (i.e., those that don't offend the reader's sensibilities). When the trials were finished, she found each of the n kids guilty, blamed each of the n kids for their bad behaviour (even if only inwardly if the reader cannot stomach the idea of blaming someone outwardly), and sent each of the n kids to their room to reflect on their behaviour, write notes of apology, and so on.

As she walked down the long corridor past the n doors to the n rooms holding the n kids, she paused before each door to think about the kid on the other side. Their evident unhappiness weighed heavily on her as she takes no pleasure in dishing out punishments. When she paused at door 1, she reviewed the evidence in her head thinks that, yes, 1 did steal 2's allowance. When she paused at door 2, she reviewed the evidence and thinks to herself, yes 2 did accidentally break 3's glasses and tried to blame that on 4. A similar pattern holds for the remaining $n-2$ cases. In each instance, it seems to Agnes that she knows that the kid in question was guilty of some specific offence. Build in what you must so that it seems to you, too, that if the relevant kid was guilty, Agnes knew it.

The plot twists. Agnes's sister, Agatha, is known by all to be all but infallible. She texted Agnes to say that precisely one of the kids was innocent. Agnes accepted that there were indeed (precisely) $n-1$ guilty kids. She knew that Agatha wouldn't say such a thing unless she knew it to be the case and waited for further texts about the identity of the innocent kid. Alas, a second text identifying the innocent kid never arrived. Agnes learns the network is down.

As things stand, Agnes is convinced in n cases that the relevant kid is guilty. She is also convinced that there are $n-1$ cases in which a kid is guilty. As she walks up and down the long corridor reviewing the cases in her mind, she's still convinced in each case that *this* kid is guilty.

There are lots of questions we might ask about this case, *The Difficult Weekend*. We might want to know what Agnes could know or rationally believe at the two stages of the story (i.e., prior to and then after Agatha's testimony). We might want to know whether she should believe kids to be guilty, blame them, and continue to punish them once she hears from Agatha that one kid (Inge, as it happens) is not guilty. We think we have answers to these questions. We'll use these answers to test theories that give us resources to try to answer questions about the way Agnes ought to respond. Our discussion focuses on a reasons-centred approach (the *reasonologist's* approach) and a decision-theoretic approach (the *expectabilist's* approach). We fear that neither account gives us the tools we need to vindicate our answers. We sketch a hybrid view that combines elements of these two approaches that vindicates our answers. While we are somewhat sceptical of a reasons-centred approach to our questions, we are ultimately open to the idea that normative reasons play an important role in our normative theorising. If reasons earn their keep, they play a very different role than envisaged by our reasonologists.

1. Should reasons be our guide?

According to the *reasonologists*, reasons should be our guide. Reasonologists, for our purposes, mean two things by this. The first is that whenever we should or should not respond in some way, this is because some reason or reasons that apply to us make it so. ('No oughts or should nots without explanatorily prior reasons to or reasons not to.') The second is that the reasons that apply to us and bear on the ways we should respond only have this bearing if they are accessible in the relevant circumstances.³

To simplify our discussion, we assume normative reasons are constituted by facts or true propositions and that these are typically (but not exclusively) facts or true propositions about the situation the agent faces (e.g., that 17 shoved 16, that n-20 put gum in Inge's hair, etc.).⁴ We see no inconsistency in saying that facts about the situation normally determine what Agnes prospectively ought to do or how she prospectively ought to feel. Dancy floats the proposal that the reasons that determine our duties (and so determine how we ought (in some sense) respond) pass through an evidence-relative epistemic filter. He remarks:

Even if duties are always grounded in features of the situation, it might be that to serve as a ground a feature

³ Readers might suspect that our reasonologists are reasons-firsters. This suspicion is correct, but the connection is more sociological than conceptual. First, a reasonologist in our sense might think that values come first and that all normative reasons are value-based. Second, a reasons-firster might think that reasons guide in the first sense but not the second. Our reasonologists think that the relevant normative reasons determine what we ought to do in the deliberative sense and that this is the important notion to theorise. They might be persuaded by the various arguments for thinking that an objective ought would not be normative in any interesting sense. A reasons-firster could think that the facts that constitute normative reasons matter primarily to an objective 'ought' and have no direct bearing on a deliberative 'ought'. A pure objectivist view or a divider's view might be out of fashion amongst the reasonologists (and might be for good reason), but it's a theoretical option that someone could defend. We should note that the accessibility of the reasons that bear on how we ought to respond might only capture part of the idea that reasons should be able to guide us. See Lord (2018) for a discussion of more demanding conditions.

⁴ This is in keeping with the work of Kieseewetter (2017) and Lord (2018) along with Alvarez (2018), Dancy (1999), and Mantel (2017). Some authors have argued that normative reasons might be falsehoods (e.g., Comesaña (2018)) or mental states, but we think (in keeping with the arguments found in Alvarez (2010), Dancy (1999), and Littlejohn and Dutant (2022)) these views are mistaken about what normative reasons are. One particularly pressing problem for the view that falsehoods might play the role of normative reasons is that we need a story about how false reasons can support (e.g., the supported options might be the ones that are best in the situations in which the false reasons are true) but (a) it is hard to see how to formulate a view that covers the case where the false reasons are inconsistent and (b) hard to see how fans of false reasons could reject the possibility that the relevant normatively potent falsehoods might be inconsistent with each other. Readers who are attracted to the idea that false reasons determine what we should do should ask whether the false reasons view provides a better treatment of our cases.

has to be one that, in some sense suitably devised, the agent is at least capable of discerning ... [F]eatures that I have no chance whatever of discerning are surely not capable of grounding duties for me. Suppose that, unknown and unknowable to me, someone has been buried alive in my garden during the night. Could this make it wrong of me to go away for a fortnight's holiday (1999: 57)?

One way to understand the filter is to think that at any given time, there's a kind of spotlight that illuminates some facts and that these illuminated facts 'seen' by the agent are the ones that work together to determine how she should or shouldn't respond. The facts that aren't illuminated are still facts and they might help to explain why there's something unfortunate or undesirable about a choice, but such facts either have a kind of merely evaluative significance or matter only to the kind of objective normativity that our targets haven't taken much interest in.

In what follows, the main things to remember are (a) that facts about how we ought to respond are explained in terms of facts about the things that the normative reasons favour, (b) that all such normative reasons are constituted by accessible facts, and (c) that such facts are typically facts about the agent's situation.

1.1 Clarifying the Challenge

Let's think about how a reasonologist might explain the following facts: Agnes ought to punish 6 and she ought to be disappointed in 6 for doing what she did. The reasonologist might say something like this:

As Agnes knew, 6 shoved 4 into the mud on the way to school. This is why Agnes should be disappointed in 6 and should send 6 to her room.

So far, so good. This fits with an ordinary way of talking about the difficult duties of being a parent. We hope we can agree that this sort of thing is the kind of thing that would warrant punishment and blame. (Feel free to make 6's behaviour worse if you need that.)

This explanation covers a very small part of *The Difficult Weekend*. How do we generalise this across the kids? For each kid that Agnes should be disappointed in, we'd have to fill in the relevant blanks:

Agnes should be disappointed in [name] and should send them to their room to think about what they've done because [fact].

The relevant fact here should be a good reason for Agnes to be disappointed with the relevant kid. In specifying the reason, we have to identify a good reason for Agnes to be disappointed in the relevant kid and so the fact must be something that a reasonable person sees as undesirable.

The facts that we plug in have to be facts (i.e., not false propositions) *and* they have to favour the relevant responses. We have to assume that our readers, like

Agnes, have an understanding of what kinds of things merit, warrant, justify, favour, etc. the relevant responses and hope that our understanding matches yours. If the facts that filled in the blanks didn't stand in the right normative relations to the relevant responses, we wouldn't explain why those responses are the ones required, justified, warranted, etc.

Let's note that we won't know when we're done filling in the details until we know who Agnes should punish, who she should blame, and who she should believe to be guilty. We think that the following is plausible:

Known Guilt: If Agnes knows that [name] is guilty, she ought to blame [name] and punish [name].⁵

The reasonologists often say that knowing or being in a position to know that p is sufficient for possessing p as a reason and take the possessed reasons to be the ones that determine how we ought to respond.⁶ If this is right, Agnes ought to blame and punish any of the n kids that she knows to be guilty. We want to note that there are additional constraints that connect types of response. We don't think, for example, that it's appropriate for Agnes to punish a kid if it's not appropriate for her to blame the kid or that it's appropriate to blame a kid if it's not appropriate for her to believe a kid to be guilty:

Believed Guilt: If Agnes ought to blame and punish [name], Agnes ought to believe that [name] is guilty.

Agnes isn't the sort of person who'd say, 'I have no idea who's responsible, but I blame you, 5!'

When the reasonologists try to describe our case, they'll have to tell us which of the n kids should be punished. Their answers must be informed by independently plausible claims about what Agnes can know, fit with what we know about the prospective 'ought', and their description should conform to the constraints that connect punishment to belief and blame.

1.2 Sharpening the Challenge

To sharpen the challenge, we'll present three arguments. Each argument purports to tell us something about how Agnes ought to respond. We think the arguments cannot each be sound because their conclusions are jointly inconsistent. We don't think that the reasonologist can reject the assumptions that generate this inconsistency. The lesson that we take from this is that the reasonologist isn't using the right tools to describe what an agent procedurally ought to respond in *The Difficult Weekend*.

Consider our first argument:

The Not-All Argument

⁵ While we think that Known Guilt is *plausible*, we think that needs to be qualified and interpreted properly. We discuss Known Guilt further in §3.

⁶ A more demanding epistemic relation could be substituted for knowledge, but we should remember that this relation should be one that we plausibly bear to facts about the situation.

P1_{na}. If there is a case in which Agnes is not aware of some fact or facts that give her sufficient reason to be disappointed in the relevant kid for doing something, to blame the kid for what she's done, this is a case in which Agnes should not be disappointed in the relevant kid for doing something or to blame the kid for what she's done.

P2_{na}. Agnes knows that there is such a case—the case in which the relevant kid broke no rule.

C_{na}. So, Agnes ought not punish all n kids, ought not blame all n kids, and ought not be disappointed in each of the n kids for what they have done.

Let's begin with a comment on (P1_{na}). There are some kinds of responses such that the absence of any good reason to respond in that way is not itself a good reason not to respond in that way. Agnes doesn't need to be aware of good reasons to whistle in order for it to be fine for her to whistle. Everyone who thinks normative reasons explain should agree that if there is no reason at all (possessed or otherwise) for Agnes not to X, it's not the case that Agnes ought not X. However, there might be some kinds of responses in some contexts where the absence of any good reason to X means that the relevant agent should not X. Blame and the expression of it might be like this. Punishment and being disappointed in someone for something they did might be, too. In each instance, the absence of a good reason to respond that way ensures that the relevant agent ought not respond in that way. (Putting this in terms friendly to the reasonologist, the absence of good reasons to X ensures the presence of decisive reason not to X.) While carrying an umbrella is not the sort of thing that would be wrong in the absence of good reasons to carry, punishing someone and intentionally inflicting harm or hardship *is* the sort of thing that's wrong in the absence of good reasons to respond in this way.

If this is right, then if (P2_{na}) is correct and Agnes knows that there's one case in which there couldn't be a fact or some facts such that she's aware of the fact or facts and this fact or facts give her sufficient or decisive reason to blame or punish, we have a good argument for the conclusion that Agnes ought not punish all and ought not blame all.

Let's consider the second argument. The conclusion of the *Not-All Argument* is compatible with the claim that Agnes ought to punish some of the kids and with the claim that she ought to punish none of the kids. We think there's an argument that the reasonologist ought to find compelling for thinking that some children must be punished:

The Some-Must Argument

P1_{sm}. If there is any case in which Agnes knows that some kid has done something unspeakably bad, this is a case in which the kid must be punished and must be blamed.

P1_{sm}. There is a case in which Agnes knows that some kid has done something unspeakably bad (e.g., Agnes knows

that 6 bullied 4, Agnes knows that 6 tried to box and ship 7's hamster, etc.).

C_{sm}. So, some kids must be punished and must be blamed.

Remember that our reasonologists take knowing to be sufficient for possessing a reason. The facts known to Agnes when she knows a kind to be guilty seem like reasonable candidates for being sufficient or decisive reasons to blame and punish. If Agnes were to refrain from blaming and punishing despite her knowledge of wrongdoing, she would systematically fail to respond to these possessed reasons to punish. It's hard to see how the reasonologist could sign off on that being appropriate given our estimation of the extent of Agnes's knowledge.

Here's our third argument:

The All-Or-None Argument

P1_{aon}. If there is some case in which Agnes ought to punish and blame and some case in which she should not punish or blame, there must be some difference between the evidence or perspective Agnes has on these two cases.

P1_{aon}. But there is no such difference between the cases.

C_{aon}. So, either each of the kids should be punished and blamed or none should be.

The first premise is the internalist constraint. When pairs of agents are internally the same, they aren't required to do different things.⁷

Our reasonologists accept this because they think denying it violates this idea that reasons provide guidance that's followable by the agents that the reasons apply to. The second seems like something we can stipulate. Agnes reached her conclusions about guilt by relying on the right kind of evidence in *n* cases. She processed that

⁷ A classic case for this is found in Cohen (1984). Our reasonologists accept this constraint. Both Kiesewetter (2017) and Lord (2018), for example, think that the same things would be rational for two individuals that are internal duplicates and think that we ought (in prospective or deliberative sense) to respond in the ways that rationality requires. That's not surprising given that they think reasons should guide us in the two senses mentioned above. In the present discussion, denying this would be akin to saying that in a preface-type case, a thinker deliberatively ought to respond differently (e.g., by believing in one case and withholding in another) despite the fact that there is nothing that the individual could point to that would explain why this differential response is appropriate. It should be noted that *if* the reasonologist thinks that a thinker's evidence includes all and everything that we know Williamson (2000) *and* thinks that in a preface-type case, it's possible to know each of the true propositions considered that there is a difference in the evidence that supports the known propositions and the one false one (i.e., the known ones are supported by entailing evidence and the false one cannot be). Few defend the view that this difference in evidential support, which seems to be unknowable to the thinker, makes a rational difference Praolini (2019). The authors above, for example, agree that evidence and reasons might not supervene upon the mental, but they nevertheless think that what's rationally required does.

evidence in just the way that every (non-sceptical) internalist would say she should. On the assumption that it's possible to have sufficient evidence to believe, blame, and punish even when the relevant party is not guilty of any wrongdoing, this is what we have in n cases.

Since the reasonologist will presumably agree that it cannot be true that Agnes ought to punish some but not all and ought to punish some only if she punishes all, they need to reject one or more of the premises in these arguments to undermine one or more of the conclusions we've just presented arguments for. We think that they cannot reasonably reject any of the premises because either they're premises that they accept or because they are things we can independently establish.

1.3 *The Reasonologists Respond*

What might the reasonologists say in response?

First Response: Some Children Must be Punished

The reasonologists might respond by rejecting the conclusion of the *All-or-None* argument. In effect, they'd say that some children should be punished and that some shouldn't be. They might say that when Agnes knows that a kid is guilty, she should punish because she's aware of a fact that gives her a decisive reason to do so. (The same would hold for believing, blaming, being disappointed with, and so on.) Thus, we arrive at the conclusion that some children must be punished. When, however, she is not aware of some fact or facts that give her reason to punish, she shouldn't punish. (In this way, punishing is to reasons to punish as belief is, on many views, to reasons to believe. Just as we shouldn't believe if there's not sufficient reason or evidence to believe, we shouldn't punish if there's not sufficient reason to do so.) Thus, we arrive at the conclusion that at least one child must not be punished. Because Inge is innocent, Agnes couldn't be aware of a fact or some facts in her case that would constitute a decisive reason to punish, to blame, to be disappointed with her, etc. So, she shouldn't punish Inge but should punish the rest.

We think the problem with this view is that it violates the idea that normative reasons should be a *guide* in the sense that it provides guidance that the agent can use. There needn't be any interesting difference in Agnes's epistemic situation when it comes to her beliefs about the guilt of 1 (who is guilty) and her beliefs about the guilt of Inge (who is innocent) beyond the fact that she knows in one case and doesn't know in the others. One way to make this clear would be to point out that on most conceptions of evidence, Agnes could have the same evidence she has in the actual world in various possible worlds where we permute the facts about the guilt of the children so that, say, Inge who is actually innocent would be guilty and 1 who is actually guilty would be the innocent one. The result would be that this view would say that even when we held all the facts about Agnes's non-factive mental states fixed and held her evidence fixed, she ought to punish 1 and refrain from punishing Inge in one set up and ought to punish Inge and ought not punish 1 in the alternative. This can't be a plausible proposal about what Agnes

prospectively ought to do. We've held fixed the facts that matter to the prospective ought.

The problem points to a problem with combining the reasonologist approach with Known Guilt. Known Guilt tells us that Agnes ought to punish and blame if she knows a child to be guilty. It says nothing about cases where this isn't known, but we know the reasonologist will say that an agent ought to respond in the same way (i.e., punish, blame, be disappointed) iff she's aware of some fact or facts that favour the same response. Our reasonologist at this point says (quite rightly, in our view) that Agnes won't be aware of a fact or facts that favour the same response if she's in a case in which the kid isn't guilty, so she insists that there's one case in which it's not true that Agnes ought to punish. Indeed, she thinks that with responses like punishment, blame, etc. the absence of any sufficient reason means that Agnes should *not* punish, blame, etc. So, just as *what Agnes knows* doesn't supervene upon facts about her evidence and her non-factive mental states, *what reasons Agnes is aware of* and *whether she's aware of a decisive reason to punish* doesn't supervene upon her evidence and her non-factive mental states. Since, however, the reasonologist is offering us an account of how Agnes *prospectively* ought to respond, this won't do because the prospectivist thinks that facts about how any agent ought to respond supervenes upon this agent's evidence and her non-factive mental states and so must also think that facts about *whether she's aware of a decisive reason to punish* must likewise supervene.⁸

Second Response: If Many Were Guilty, All Should be Punished

Let's consider an alternative line of response. The reasonologist might reject the conclusion of the *The Not-All Argument*. In other words, they might argue that each of the children should be punished. They might say that since Agnes knows that 1 is guilty, she must be aware of some fact or facts that give her decisive reason to punish. In keeping with the goal of telling us how Agnes *prospectively* ought to respond, they'll say that Agnes ought to treat the remaining kids the same.

If the reasonologist pursues this line, they'll have to say that in each case, Agnes is aware of some fact or facts constitute normative reasons that give sufficient or decisive reason for blaming, punishing, being disappointed with, etc. In n-1 cases, it's not hard to identify these reasons—they'll be facts about deception, about

⁸ Of course, the reasonologist might then say that they'll adopt a different conception of evidence, one on which Agnes's evidence differs in the scenarios described *because* the facts about guilt differ. (They might adopt E=K, for example. See Williamson (2000).) The trouble here is that this view of evidence isn't available to the reasonologist if they're giving a prospectivist view of how Agnes ought to respond. Dialectically, this commits them to views about what Agnes ought to believe according to which the beliefs that Agnes ought to have, say, in the good case and in the bad case would differ. We don't think it's plausible for a prospectivist to say that their view is true to the prospectivist idea if it also implies that, say, Agnes should believe she has hands if she's in the good case and shouldn't believe that if she's a BIV.

bullying, about theft, about the mistreatment of animals, etc. It's the nth case that's the difficult one. Inge, recall, was innocent. It's not clear what fact or facts Agnes could have been aware of in this case that we'd propose to fill in these blanks:

- (i) Agnes's reason for being disappointed in Inge was a good one, namely that [fill in the blank].
- (ii) There was a fact that Agnes was aware of that gave her a decisive reason to punish Inge, which was that [fill in the blank].

In other settings, reasonologists have proposed that back-up reasons can do the normative work that would have been done by the facts about the situation that our misinformed agents would cite in trying to defend or justify their responses.⁹

What should we think of this thesis?

Surrogate: An agent is aware of some fact or facts that give her decisive reason to X iff she is the non-factive mental duplicate of some agent who is aware of some fact or facts that give her decisive reason to X.

If the Surrogate Thesis is correct, Agnes would be aware of some fact or facts that give her decisive reason to punish Inge and decisive reason to be disappointed with Inge if she is the non-factive mental duplicate of some agent who is aware of some fact or facts that give her decisive reason to punish, be disappointed with, etc. We're assuming that *if* (contrary to fact) Inge had been guilty, Agnes could have known this, the Surrogate Thesis tells us that there must actually be some fact or facts Agnes knew that gave her decisive reason to punish, blame, be disappointed with, etc. even though Inge wasn't guilty of anything. What might this surrogate or backup reason be? It could be that Inge looked guilty (or the facts in light of which she looked guilty). We don't know the identity of the fact, only that there must be facts that are distinct from those in virtue of which Inge would have been guilty that favour the same kinds of responses as the facts in virtue of which Inge would have been guilty. These reasons, in effect, would do all the things that the reasons constituted by Inge's guilt could have done.¹⁰

⁹ See Kiesewetter (2017) and Lord (2018).

¹⁰ Readers should feel free to check this for themselves. By hypothesis, Inge's behaviour was impeccable, but things did not seem that way to Agnes. We think there is no fact known to Agnes that would be a good reason to be disappointed in Inge. Such reasons would have to point to things that are not just undesirable or unwanted but things about Inge's conduct that were unwanted or undesirable (e.g., to us, the fact that Agnes mistakenly but rationally believes Inge did something wrong helps explain why she's rational in blaming her and being disappointed in her, but it is not itself among the good reasons there might be for being disappointed in someone; if reasonologists disagree, we disagree about the support or favouring relation that holds for attitudinal responses like disappointment). We of course agree with the reasonologist that in the bad case, a thinker might rationally be disappointed in someone who did nothing wrong or that the thinker ought (in the prospective sense) be disappointed in them, but we find no known fact that

We're sceptical that there are these facts that constitute decisive reason to be disappointed with Inge, punish Inge, etc. that would be known to Agnes even when Inge is innocent. Remember that the reason must meet three conditions: it must be a fact, it must be possessed by Agnes, and it must play the right normative role. Suppose Agnes believes Inge bullied n-8 and expresses this in explaining why she's punishing Inge. The following candidate facts seem unpromising as good reasons for this response:

- That Inge bullied n-8;
- That Inge probably did something wrong/probably bullied n-8;
- That most of the kids did something bad;
- That someone said Inge bullied n-8;

The first candidate isn't a fact. The second *might* be something Agnes knows, but it's not a good reason for Agnes to be disappointed in Inge. The same holds for the third and fourth. We couldn't think of a better fifth candidate. Remember that the facts we are looking for are, *inter alia*, facts that would be good reasons for Agnes to be disappointed in Inge. Such reasons, we think, have to be facts that reveal (not merely evidence or indicate) that Inge's behaviour was undesirable in some way. Inge's responsibilities, we hope reasonable people can agree, might be to treat her siblings well, but would not include controlling the evidential inputs that Agnes has to go on in updating her beliefs. Agnes isn't Caesar. Agnes might be disappointed in one of the kids for behaving badly but not for being under suspicion.

We think that the availability of backup reasons in error cases should turn on *substantive* commitments about what kinds of things might justify or require certain kinds of responses. The Surrogate Thesis is not defended on any substantive grounds. We don't see why we should agree, without any very surprising substantive argument about the kinds of things that count in favour, that because there's some possible world in which Agnes knows Inge bullied n-8, in every possible world Agnes knows some other fact that equally gives Agnes a good reason to be disappointed in Inge. (What if Inge were asleep? What if Inge were dead? What could these facts be?) We do not accept this conditional: if Agnes were aware of some fact or facts that in 2's case constituted a good reason for being upset with 2, in any case in which Agnes is in the same non-factive mental states, she would also be aware of some fact or facts that constituted a good reason for blame or punishment. When we think about the kinds of facts that would supervene upon Agnes's non-factive mental states, in some of the cases, they would include *only* facts about her own mental life, how things seemed to her, what her evidence supported, etc. Such facts are *not* typically good reasons for being upset with someone. The kinds of facts that could both be a subject's reason for being upset with someone and a good reason for being so upset have to be (a) things that the subject believes to obtain and (b) things that the subject wants not to obtain. The fact that, say, Inge broke a rule would meet (a) and (b) but there is no such fact. The fact that it *seemed*

constitutes a good reason for this response and so, predictably, question the utility of a reasons-centred explanation of the intuitive data about the case.

that Inge broke a rule might satisfy (a) but not (b). Agnes's reason for being upset, if she's reasonable, wouldn't be a fact about how things seemed to her because her relevant wants would have to be about Inge and her behaviour, not Agnes and how things struck her.

Alvarez (2018) illustrates the point well. Othello is upset with Desdemona. He believes that she was unfaithful. If we changed the details of the play, Othello could have come to know that she was unfaithful from Iago's testimony. In this altered version of the story, Othello's reason for being upset with Desdemona would be that Desdemona was unfaithful. Perhaps this fact would be a good reason for him to be upset. (We can bracket the other aspects of the play because they are not relevant to the point about the reasons that support emotional responses.) Compare this retelling of the story to the play as written. In this case, Othello did not know that Desdemona was unfaithful because she was faithful. Iago lied. According to the backup reason proposal, since there is no difference in Othello's non-factive mental states, it's just as rational for Othello to be upset with Desdemona in the two versions of the story. And according to the idea that rationality consists in responding correctly to reasons, reasons perspectivists will say that in each case, Othello is rational in being upset with Desdemona only if in each case he knows some fact or facts that constitutes a sufficiently good reason for being upset with Desdemona. The backup reasons view says that *if* Desdemona's infidelity was a sufficient reason for Othello to be upset, the fact that it appeared to Othello that Desdemona was unfaithful would also have to be a sufficient or decisive reason for Othello to be upset with her.

This cannot be right. We do not think that Othello's reason for being upset could have been that it appeared to him that Desdemona was unfaithful or some facts that don't include facts about infidelity (e.g., facts about how things appeared, about Othello's evidence, etc.).¹¹ And we do not think that he would think that this reason would be a *good* reason for him to be upset. It is no part of Shakespeare's play that Othello would have thought that a good reason for being upset could have been that things appeared or looked some way. This is inconsistent with his values. There are characters who would be upset by the fact that it appeared that someone was unfaithful (e.g., people who would be upset with someone because they didn't control how things would look to others). We can see why Shakespeare would have chosen not to tell his tale with such a character as the centre of the action. Othello's reasons would have to do with fidelity, not the appearance of it, because he valued fidelity, not the appearances of it.¹² We should not generally assume that if the fact

¹¹ On this point, we are in agreement with the defenders of the false reasons view.

¹² One way to put it might be this. An agent's reason for being upset would be something that the agent finds upsetting or something for the agent to be upset about, not something that merely explains why the agent is upset. Othello does not think that the fact that it appears that Desdemona is unfaithful is something to be upset with her about. Here it is important to mark the difference between explanatory reasons and motivating reasons.

that p constitutes a decisive reason for an agent who knows p to Y , the fact that it appears to an agent that p constitutes a sufficiently good reason for some mentally similar agent who knows this fact to Y . As the case of emotion illustrates, some agents who would agree that p is a reason for an affective response would deny that the appearance that p would be any reason for that response.

Third response: Free the Children

Reasonologists might reject the conclusion of the *Some-Must Argument*. Given that they're offering a prospectivist view, it seems that each case should be treated the same. Since she knows there's one kid she shouldn't punish, couldn't Agnes know that she shouldn't punish any?¹³

The disagreement we have with this reasonologist is, largely, a disagreement about the epistemic side of things. We can both agree that Agnes will possess the fact that a kid is guilty as a reason if she knows it, but the reasonologist seems to think that even if Agnes knows initially that the guilty kids are guilty, she cannot rationally believe them to be guilty in light of the sister's testimony. It seems to function as a defeater and if Agnes doesn't rationally believe p at the end of *The Difficult Weekend*, she doesn't possess p as a reason at the end of the *The Difficult Weekend*.

We think that this line is not convincing. On the one hand, it seems that, upon pain of scepticism, they must concede that she could have known prior to the testimony that the guilty kids were guilty and so everything is riding on this appeal to defeaters. We think this appeal to defeaters is unpersuasive.

We often know that a potential source of information (e.g., a phone book, a set of encyclopaedias) contains an error that hasn't been identified and yet individual claims from this source seem to be very probable candidates for knowledge. We should add that we're convinced by Christensen's (2004) suggestion that many of us have inconsistent beliefs where we cannot identify the false belief where it seems the beliefs in question are paradigmatic cases of knowledge. If knowledge is sufficiently robust to persist in the face of this recognition, we shouldn't be moved by an appeal to defeaters.

We think that it's helpful to think about what happens when someone like Agnes commits to memory, say, the entries in a set of encyclopaedias and we ask how many of these individual beliefs might constitute knowledge. We think that, in principle, it could include all the beliefs corresponding to true entries. In learning that one entry is mistaken and knowing that she cannot identify the least supported

¹³ This is the line that we would expect someone like Smith (2018) to take. He thinks that in preface-type situations, we cannot have evidence that provides sufficient support for each proposition in a set known to be inconsistent Smith (forthcoming). If, as we've been assuming, the support for each is essentially the same in terms of their epistemically relevant properties, he thinks we lack justification to believe in each case. And we think, in keeping with his outlook, that if we shouldn't believe someone to be guilty, we shouldn't blame and shouldn't punish.

member of this set, we think that Agnes could retain all the knowledge she had initially. If she retains this knowledge, however, it seems that she should be able to retain her rational beliefs corresponding to the entries since clear cases of knowledge look like clear cases of rational belief.¹⁴

We think that the defeatist response leads to scepticism. Think about minor variations on the preface cases we've seen where the belief about the fallibility of the source is based on testimony. Compare the case in which Agnes has committed n beliefs to memory from an encyclopaedia she learns contains $n-1$ true entries to a case in which, say, Agnes* has committed n beliefs to memory from an encyclopaedia that she learns has a .1 chance of containing 11 errors (and .9 chance of being error free) or a .05 chance of containing 22 errors (and .95 chance of being error free). In these setups, the expected amount of ignorance or mistaken belief is greater than it would be in the case in which Agnes is certain there's one mistake in n cases. If we say that in these chancy setups, Agnes* must suspend because her situation is worse than it is in the original, we're going to get dangerously close to scepticism. If we say instead that she needn't suspend in these cases, since these cases seem worse than the case in which Agnes is certain there's one error in n cases, we shouldn't say that Agnes ought to suspend in this case. We see no reason to think we ought to prefer the policy of suspending to believing in the less desirable setup if believing is permitted in the worse setup that Agnes* faces.

1.4 Summing Up

We do not believe that there is any set of facts that seem like good reasons to punish or be disappointed with one of the kids that might directly determine how Agnes ought to respond. The reasonologists cannot appeal to unknown facts in their explanations because they're trying to explain how Agnes *prospectively* ought to respond. They cannot appeal to the known facts to explain the relevant normative data because there is not a one-to-one correspondence between the cases where the known facts favour punishment or favour freeing a kid and the cases where Agnes ought to punish or ought to free a kid. As Agnes knows, there is one case in which she ought to punish where she could not have been aware of a fact that constituted a good reason to punish or be disappointed with the relevant kid. Alas, she knows not which case this is.

In some ways, the difficulty isn't surprising. Possessed reasons are taken to be things the agent knows or bears some knowledge-like relation to. Because reasonologists must concede that facts about which reasons the agent possesses will not supervene upon the agent's non-factive mental states and relevant evidence, we should doubt that the reasons as reasonologists understand them will explain the normative data since the facts about the ways an agent ought to respond will supervene upon the agent's non-factive mental states and relevant evidence. Once we see that the sets of possessed reasons don't supervene upon the things that

¹⁴ We spell out some of the bad consequences of denying this in Dutant and Littlejohn (forthcoming).

determine how the agent prospectively ought to respond, we should worry about any attempt to explain the latter in terms of the former.

2. Aren't probabilities the very guide of life?

Here's a natural thought. The thing or things that explain why Agnes prospectively should or shouldn't respond in some way should supervene upon her non-factive mental states and relevant evidence. Because the things she knows and the reasons she possesses don't so supervene, we should look for something else to explain why Agnes should or shouldn't do this or that. A natural alternative might be the agent's rational degrees of belief. If probabilities are the very guide of life, maybe reasons won't be.

Let's consider an alternative to the reasonologist's view, expectabilism.¹⁵ According to one version of this view, A ought to X iff (and because) X-ing maximises expected value.¹⁶ On this view, we might be able to say some things to accommodate reasons-talk (e.g., that Agnes has 'most reason' to do such and such because the expected value of doing such and such is greater than that of doing anything else instead), but it's clear that reasons themselves are not doing much work—the real explanation as to why Agnes ought to punish 1 is that there's something desirable about punishing in some circumstances, something undesirable about punishing in others, and once we know the magnitudes of these values and the probabilities of these possibilities, we can rank our options and choose accordingly. If facts about rational degrees of belief and the values of outcomes supervene upon an agent's non-factive mental states, we should expect that the view will deliver the kinds of internalist-friendly verdicts that fit with the prospectivist outlook.

Here's how an expectabilist might handle questions about when it's right to punish. When Agnes is faced with a decision about whether to punish, she knows that her options are to punish or refrain and that the child before her is either guilty or innocent. In each case, knowing that n-1 kids are guilty, there's some uncertainty as to whether the kid is guilty. As she's equally confident in each case that the kid is guilty, we'll assume that the probability of guilt in each case is n-1/n.

	Guilty	Not Guilty
Punish	A	b

¹⁵ We find expectabilist approaches to questions about the permissibility of punishment in a number of places, particularly in discussions of the proof paradoxes. See Papineau (2021), Ross (2021), and Steele (ms.).

¹⁶ One worry about expectabilism, as understood here, is that value does little to no work in cases similar to ours. When it comes to belief, some of us share Raz's scepticism about value or value-based reasons explaining why we should have attitudes like belief or blame (Raz 2009). We also think that the operative reasons in our cases will involve rights and duties that cannot, in turn, be understood in terms of values for reasons highlighted by Kieseewetter (forthcoming).

Refrain	C	d
---------	---	---

Agnes thinks about the possible outcomes (a-d), knows their values, and so can determine the expected value of punishing and refraining. Provided that the values we assign to the potential outcomes are sensible and n is sufficiently large, Punish maximises expected value in each case. So, the expectabilist can explain why Agnes ought to punish each of the kids.

The problem with this proposal is that it gets this case wrong:

The Garden Party

n children are playing in the garden when suddenly $n-1$ children put into action their complicated plan for tormenting the neighbour's cats. The other child knew nothing about this plan and saw nothing as she was busily reading, or colouring, or doing whatever it is that good children do. There is no further information that we can use to settle the question of any particular child's involvement.

In this case, we think that Agnes cannot blame and cannot punish any of the kids even though it's highly probable in each case that the kid is guilty.¹⁷ If the expectabilist case for Punish in *The Difficult Weekend* is sound, an expectabilist argument for Punish in *The Garden Party* should be, too. The cases don't differ in a way that should matter to the expectabilist.¹⁸

3. What should be our guide?

Neither the reasonologists nor the expectabilists get our cases right. The reasonologists cannot vindicate our intuitions about the *The Difficult Weekend* (i.e., that Agnes ought to punish each of the kids despite knowing one is innocent). The expectabilist gets this case right, but they bungle *The Garden Party* (i.e., that Agnes ought not punish any of the kids even when each is quite likely to be guilty). Here is our diagnosis. We need resources from both these views. We need the expectabilist's probabilities to handle *The Difficult Weekend*. What we learn from *The*

¹⁷ We side with Adler (2002), Buchak (2013), Littlejohn (2020), Smith (2018), Thomson (1986), and Moss (2018) in thinking that naked statistical evidence isn't sufficient to warrant blame or punishment in these cases. We think it also doesn't suffice to warrant (full) belief.

¹⁸ We also think that this kind of case poses problems for moral encroachment views if their explanation of our intuitions about *The Garden Party* is given in terms of a threshold that shifts or changes because of the (salient) moral stakes. On this view, it's rational to believe (and blame, perhaps) when there's sufficient evidence to believe and the (salient) moral stakes can help to determine when the strength of support is sufficient. If this were the right approach, we'd expect there to be similar thresholds in these cases (since the salient moral stakes are the same). This approach would seem to predict, then, that it's permissible to believe someone to be guilty in one case iff it is in the other. But this *isn't* what we find.

Garden Party is that we need some guidance to figure out which probabilities matter (e.g., the probability of 9's guilt might be the same in both cases, but blame seems inappropriate in one case and not the other). We know that the probability of guilt isn't the thing that matters to blame, as evidenced by our reaction to *The Garden Party*. We won't know which probabilities matter in answering questions about the ways Agnes should or should not respond until we know substantively what matters to the objective suitability or unsuitability of our responses. It is here that reasons might earn their keep. We think normative reasons might tell us, in the first instance, which responses would be objectively suitable or unsuitable. Possible sets of reasons and the probabilities, in turn, determine the ways Agnes prospectively should respond. Actual sets of reasons determine how she objectively should respond. In the normal course of things, the notion of the possessed reason seems to have no interesting role to play as possession would typically matter little to the objective side of normativity and binary epistemic notions like knowledge seem to matter little the prospective side of things in the normal course of things.

How might we combine the reasonologist and expectabilist views? We identify the facts that matter to objective suitability, rightness, etc. We assume that such facts favour or disfavour certain responses and might (sometimes) do so without taking account of the agent's epistemic circumstances. We also assume that some such considerations are weightier than others. Other things equal, it is worse to act against a weightier reason, for example. Our theory of normative reasons informs our account of what factors matter to objective suitability, rightness, etc. and how undesirable or wrongful it might be to respond in ways that go against the reasons.

Actual sets of such reasons and their weights will determine how objectively wrongful various responses might be. We assume that nobody has a good view about how reasons combine, but we assume (setting aside complications arising from supererogation or justifying reasons that permit us to respond in suboptimal ways) that objective reasons give us a way to rank things from the most to least wrongful and that our responses objectively ought to be one of the least wrongful. The probabilities that 'matter' concern possible sets of objective reasons. The expected objective wrongfulness of an option is determined by (a) how wrongful the response might be relative to these different possible sets of reasons and (b) how probable it is that some such set of reasons is actual. We assume that we generally prospectively ought to minimise expected objective wrongfulness.

This proposal cannot be applied to the cases at hand without importing substantive assumptions about what makes responses like blame or punishment objectively suitable. In general, we think that some of the facts that obtain can have a normative bearing regardless of the epistemic relations we bear to them. Inge's innocence, for example, might have made it the case that Agnes should not blame her even though this fact is not known to her, not reasonably believed by her, is not

part of her evidence, and is a possibility she rationally takes to be quite remote. Some of the facts, however, that seem to have a normative bearing might have an epistemic dimension but nevertheless will bear on what's objectively suitable even if we bear no epistemically interesting relation to this epistemically loaded fact. That Agnes knows 9 to be guilty, for example, might be the one thing that makes it the case that she ought to punish 9 and blame her. And it might do that even if Agnes doesn't know that she knows. Our cases, we think are interesting and instructive, because the facts that we think must obtain for it to be the case that Agnes objectively ought to punish or blame have an epistemic dimension and related non-epistemic facts do not favour the same response. The bare fact that 9 is guilty, for example, is not something that makes blame objectively suitable, on our view. The epistemically loaded fact that 9's guilt is known, however, is. The bare fact of 9's guilt always either obtains along with the fact that it's not known (in which case, blame is objectively unsuitable) or with the fact that it is known (in which case, blame is objectively suitable but the bare fact adds nothing). We think blame is interesting here because unlike, say, making bets or kicking rocks, the possible obtaining of the desirable result constitutively depends upon an epistemic relation to a fact where this fact, on its own, does not make the result desirable. The prize, as it were, requires knowledge whereas in taking a bet, knowledge means that a prize is within reach where this prize (e.g., a monetary reward) and its desirability does not constitutively depend upon this knowledge.

When it comes to believing, blaming, and punishing, it matters that the subject can identify good reasons to blame, to believe, and to punish. Knowledge should suffice for the possession of such reasons, but we also think that knowledge is necessary for the possession of the relevant reasons in this case. When it comes to punishment and blame, for example, there are *crucial* reasons, reasons that a subject must have in order to properly punish or blame:

When combined, doesn't this lead right back to the problem that Agnes ought to punish some but only some of the kids? These ideas imply that Agnes should punish some and spare others (assuming Agnes has the knowledge we credit to her) but why is this a problem? It's only a problem if Known Guilt and Only Known Guilt are taken to be claims about what Agnes prospectively ought to do, but what if we read them instead as claims about what she objectively ought to do? Our proposal is that awareness of guilt is necessary and sufficient for blame and for punishment to be objectively suitable.

We think that if normative reasons are taken to be facts about the situation, they determine whether Agnes objectively should or shouldn't respond in some way. As a substantive matter, we think that the facts that matter to objective rightness are facts about what Agnes knows or what Agnes is aware of (and not the facts known or that she's aware of).¹⁹ If normative reasons and their weights make

¹⁹ For a discussion and defence, see Littlejohn (2020).

some part of the normative world go round, they matter in the first instance to the objective dimension of the normative world.²⁰

On our view, the reasonologists made a subtle mistake by proposing that awareness of facts matters in the ways they've suggested to prospective rightness. First, while we think awareness of certain facts is necessary for right punishment and right blame, we think this is a special case, a case that differs from cases like betting on a football match, deciding to take an umbrella, or pouring the contents of this bottle into a friend's glass. We don't think in these more mundane cases, the permission to act depends upon whether there's some fact or facts that the agent is aware of that favour these actions. The case of punishment (or, for that matter, praise) might be somewhat different from the mundane cases because of the expressive nature of these acts (i.e., that they are supposed to express attitudes towards certain states of affairs). The propriety of these acts requires the agent to be suitably related to the wrongdoing being condemned (or the laudable behaviour being praised), but we don't think this pattern holds in the mundane cases. Second, we think, as the expectabilists do, that binary relations between agents and actual facts aren't really the sort of thing that matters to determine prospective rightness. Surely risk must also matter here. Once we take account of the risks by taking account of the agent's rational degrees of belief, we don't think the actual presence of the facts the agent would offer in defence of her behaviour have to obtain.

On our way of thinking about punishment, an agent like Agnes faces two possible outcomes that would be undesirable. One undesirable outcome would be to fail to punish someone she could know to be guilty given the evidence before her. Failing to take account of her epistemic position could mean she's guilty of a kind of wrong of omission. Another would be to punish someone she didn't know to be guilty. She would punish without the awareness she needs for that punishment to be objectively suitable because she'd lack the awareness she'd need for blame to be objectively suitable. In light of these risks, what should she do? We think she prospectively should respond in ways that minimise expected objective wrongfulness.²¹

A natural thought, in keeping with the spirit of Blackstone's remarks, is that it would be worse for Agnes to violate Only Known Guilt and to punish someone without awareness of this person's guilt than to violate Known Guilt and fail to punish someone she could know to be guilty given her information. If this is right, our proposal implies that it would be prospectively wrong for her to punish unless

²⁰ This preserves part of the reasonologist view. We don't have to deny that normative reasons provide guidance in the sense that they determine what (in some sense) an agent ought to do, but if they determine what the agent *objectively* ought to do (and don't determine what the agent prospectively ought to do), we don't think that they provide guidance in the second sense that matters to the reasonologist. They might not establish guidelines that we can follow owing to ignorance, uncertainty, or mistake.

²¹ Although their views differ in important ways, the views of Lazar (forthcoming) and Olsen (2018) deliver the same verdicts in our cases.

it's more likely than not that she can know that the relevant kid is guilty. So long as the difference in the weight we assign to the reasons to conform to Known Guilt and Only Known Guilt isn't extreme, we get the result that the probability that she's in a position to know a kid to be guilty might be significantly less than 1. Assuming, as we have been, that Agnes rationally assigns roughly the same probability to the hypotheses that 1 is the innocent one, that 2 is the innocent one, ..., that n-1 is the innocent one, and that n is the innocent one, we can say that she prospectively ought to punish each in *The Difficult Weekend* since the probability she's in a position to know in each case is quite close to 1. In *The Garden Party*, by contrast, the probability she's in a position to know is 0 and so our view is that it's objectively and prospectively wrong for her to punish and blame since it's certain she'll violate Only Known Guilt.

Our view differs from the expectabilist view above because our first expectabilist didn't pay any mind to the factors that determine objective suitability.²² But if there is such a thing as objective rightness, we cannot imagine any reason to ignore information about the presence or absence of objective right-making features in formulating a theory of what we prospectively ought to do. Given that a rational agent can be uncertain about whether such features are present or absent and that the degree of uncertainty helps to determine whether someone should or shouldn't take a risk, it seems that the expectabilist ought to say that the credences that figure in the theory of what we prospectively ought to do can include credences about objective right making features. This is what was missing from the first expectabilist view — credences about the presence or absence of objective reasons (e.g., those to conform to Known Guilt and to Only Known Guilt).

We claim that substantively speaking, it's wrong to violate our two norms, so the facts that determine whether Agnes conforms to these norms are the normative reasons that matter to the objective ought and her credences about these

²² In some ways our complaint about the first expectabilist view is similar to Buchak's (2013). It differs, however, in a few respects. We think the expectabilist should take account of the grounds that determine objective rightness and that this might be something other than a value (e.g., we think it sometimes has to do with reasons to conform to norms). We agree that facts about what the subject can know and should believe matter to the permissibility of punishment, but for us that's because knowledge figures in our account of when punishment is objectively suitable. It's not part of our story that belief matters when deontological constraints are in play. We haven't discussed this, but we think such constraints can be in play and it might be right to act knowing there's a risk of violating such a constraint even when it's not right to form beliefs about whether this constraint has been satisfied or violated. Our proposal explains the role of belief and knowledge in terms of the connection between belief and blame and punishment's expressive function. If we combined our proposal about when it's objectively right to punish (i.e., when we would conform to Known Guilt and to Only Known Guilt) with Buchak's, it seems that we'd be saddled with the view that it would only be right to punish when the agent can rationally believe that she knows. We don't accept that. We think having sufficiently high confidence that the subject is in a position to know should be enough.

reasons matter to the prospective ought. On the expectabilist view we're envisaging, Agnes objectively ought to punish if she knows and objectively shouldn't punish if she doesn't know. She prospectively ought to punish if it's rational for her to be sufficiently confident that she's in a position to know and prospectively shouldn't punish otherwise.

To get the case right, we think we need to borrow something from the reasonologists, borrow something from the expectabilist, and (importantly) we need to use the right tools for the right jobs. Our expectabilist was right to use credences in the theory of prospective suitability, but they didn't have a theory of objective suitability to draw from, so normative reasons to conform to norms like Known Guilt and Only Known Guilt fell out of the picture. We think reasons to conform to norms like Known Guilt and Only Known Guilt matter to objective suitability, so ignoring information about these reasons explains why their view isn't quite right. The reasonologists gave us the tools we needed for the theory of objective suitability, but they're making a mistake if they use those tools for anything but the theory of objective suitability. Sets of reason-constituting facts aren't the right tool for giving a theory of prospective suitability. Imposing a binary epistemic relation between an agent and these reason-constituting facts doesn't give us what we need because it doesn't give us what we need to explain how to deal with uncertainty.

Our positive proposal for dealing with cases like *The Difficult Weekend* and *The Garden Party* involves three parts.

- I: Normative reasons and their weights determine which responses would be objectively right.
- II. When an agent is in a situation in which it's possible that there are different sets of normative reasons that bear on what's objectively right, the agent's credences about such reasons (and their absence) helps determine which responses would be prospectively right.
- III. In the cases of belief, blame, and punishment, the relevant reasons that determine objective rightness are facts about what the relevant agent can or cannot know.

We see the first two parts of our proposal as perfectly general suggestions about objective and prospective normativity. We don't see how we can approach our cases without some substantive assumptions about when belief, blame, and punishment would be objectively suitable, so this is where our third assumption comes into play. When we combine them, we can say that in *The Difficult Weekend*, belief, blame, and punishment are appropriate because it's nearly certain (given Agnes's information) that she knows in each of the cases. In *The Garden Party*, belief, blame, and punishment are inappropriate because it's certain (given Agnes's information) that she cannot know in each of the cases.

And why, the reader might wonder, is it that facts about Agnes's knowledge matter here? We haven't said much about this substantive assumption yet, but the rough idea is this. We think that a thinker's beliefs are supposed to provide that

thinker with facts that can guide their reasoning. We think beliefs do this iff they constitute knowledge. Thus, facts about what a thinker can or cannot know about some matter are the ones that determine the objective suitability of her beliefs. When it comes to blame and punishment, we think that the ability to identify factors that show someone to be worthy or blame and worthy of punishment matter to the objective suitability of these responses. In the case of blame, these factors are what we cite in blaming. In the case of punishment, these factors matter insofar as punishment must be, *inter alia*, something suitable to express blame. The risk that matters to prospective rightness here is the risk of blaming or punishing without being aware of factors that make someone worthy of being blamed and punished. This is why knowledge figures in our substantive proposal about what makes blame and punishment objectively suitable.

We think punishment (and praise) are special cases. We don't think it's generally true that objective rightness requires knowledge of much of anything, so we don't think it's generally true that the possession of normative reasons is required for objective suitability (e.g., when it comes to the objective suitability of not backing over your neighbour, not driving backwards while your neighbour stands behind your car is objectively suitable).

Similarly, we don't think that knowledge or the possession of normative reasons generally matters to prospective rightness. Consider the much-discussed miner's case. There are ten miners trapped below ground in one of two mines and you have no evidence that indicates that it's more likely that they are working in one shaft rather than the other. It starts to rain. If you do nothing, they'll each drown. If you block one shaft completely, the other will flood completely and so you'll either guess correctly and save all ten or guess incorrectly and ten will be lost. If you partially block the shafts, that guarantees enough air in both shafts to ensure that nine will live regardless of which shaft they're in. The standard response seems to be that we should partially block both shafts. While the expectabilist explanation of this seems to us quite illuminating (i.e., partially blocking the shafts maximises expected value or minimises expected objective wrongfulness), we don't see possessed reasons as doing much explanatory work in cases like this. To be sure, the reasonologists might say that among the possessed reasons might be the fact that the act in question will maximise expected value, but it's not clear why *knowledge* of such a fact is required to explain the relevant data. Here it seems that the fact that the act in question, say, maximises expected value explains why the agent ought to partially block the shafts.

Conclusion

In much of the recent work emphasising the importance of reasons to understanding normative notions, reasonologists have proposed reasons-theoretic accounts of normative notions and of rationality. For the most part, the main concern has been to give a reasons-centred account of the prospective ought or of rationality. To do this while recognising reasons as grounded in features of a mind-independent

reality, it seemed an epistemic condition needed to be placed on the reasons. By placing such a constraint, a circumscribed set of facts could be identified as the reasons that have normative bearing. On such an approach, it is not clear what normative bearing features of the situation unknown might play. It is not clear what normative bearing facts about our uncertainties about these facts might play. Eager to find a normative role for the latter, it might be tempting to dispense with reasons entirely, but this might be premature. We have suggested that reasons might play a role in characterising the *objective* side of normativity. There might be some areas where knowledge or awareness of some fact matters to objective suitability or rightness, but we don't think that the notion of a possessed reason generally does much work in understanding this part of the normative domain. When it comes to understanding the prospective side of things, we think some nexus between the objective and prospective must exist and that it is here that probabilities matter. Expected objective wrongfulness can be used to order options, say, and the notion of a possessed reason no role in characterising this notion if it plays no role in understanding either expectation or objective wrongfulness.

Where does this leave us in terms of thinking about future work in this area? Let us highlight two things of importance. First, we have given no reason to think that reasons matter and no reason to think that they do not. We have argued that they do not play a certain theoretical role and suggested that if they matter, they might matter to objective rightness or suitability. There are considerable doubts about whether there is anything of normative significance that floats free from the information available to agents whose responses we might evaluate. A proper defence of continued interest in reasons might require two sorts of argument. We need arguments to think that reasons are important to understanding some normative notions such as arguments that we need reasons to understand rights and correlative duties (Kiesewetter (forthcoming)). We need further arguments to think that such things might be a part of a normative domain that is objective in the relevant sense (Graham (2010); Littlejohn (2012)).

Additionally, we face some very difficult questions about how to pursue the kind of hybrid view we envisage. Recall Dreier's conjecture that every plausible moral theory can be consequentialised (Dreier 2011: 97). We have doubts that every plausible *normative* theory can be consequentialised. In particular, we doubt that we can consequentialise the norms that govern responses like blame and belief. When we believe and blame as we should, we doubt that we can be represented as if we are maximising some sort of value in responding some way rather than some alternative way (Berker (2013); Elstein and Jenkins (2020)). It is not obvious what rules or principles determine which response out of some set of alternative responses is most desirable if the alternatives cannot be ordered or this ordering does not determine which response is appropriate. We can gesture at some vague ideas that work well enough for our purposes here (e.g., that the force of the reasons to believe or refrain from believing what we wouldn't can be strengthened or diminished depending upon how likely it is that our beliefs would be or wouldn't be

knowledge), but we shouldn't pretend that this is a general recipe for understanding how we ought to respond to potential sets of normative reasons given our uncertainty about which reasons actually obtain. Our hope is that future work in this area involves a collaborative effort between the philosophers who see reasons as a useful theoretical tool and the philosophers who think that we should use decision-theoretic machinery to address questions about the prospective ought.²³

References

- Adler, J. (2002), *Belief's Own Ethics* (Cambridge, MA.: MIT University Press).
- Alvarez, M. (2010), *Kinds of Reasons* (Oxford: Oxford University Press).
- Alvarez, M. (2018) 'Reasons for Action, Acting for Reasons, and Rationality', *Synthese* 195/8: 3293-310.
- Berker, S. (2013), 'Epistemic Teleology and the Separateness of Propositions', *Philosophical Review* 122/3: 337–93.
- Buchak, L. (2013), 'Belief, Credence, and Norms', *Philosophical Studies* 169/2: 1–27.
- Cohen, S. (1984), 'Justification and Truth', *Philosophical Studies* 46/3: 279-95.
- Comesaña, J. (2020), 'A Plea for Falsehoods', *Philosophy and Phenomenological Research* 100/2: 247-276.
- Dancy, J. (1999), *Practical Reality* (Oxford: Oxford University Press).
- Dreier, J. (2011), 'In Defense of Consequentializing', in M. Timmons (ed.), *Oxford Studies in Normative Ethics* 1 (Oxford: Oxford University Press), 97–119.
- Dutant, J. and Littlejohn, C. (forthcoming), 'What Is Rational Belief?' *Noûs*.
- Elstein, D. Y. and Jenkins, C. S. I. (2020), 'The Truth Fairy and the Indirect Epistemic Consequentialist', in P. Graham and N. Jang Lee Linding Pedersen (eds.), *Epistemic Entitlement* (Oxford: Oxford University Press), 344–60.
- Graham, P. A. (2010), 'In Defense of Objectivism about Moral Obligation', *Ethics* 121/1: 88–115.
- Graham, P. A. (2021), *Subjective versus Objective Moral Wrongness* (Cambridge: Cambridge University Press).
- Kiesewetter, B. (2017), *The Normativity of Rationality* (Oxford: Oxford University Press).
- Kiesewetter, B. (2018), 'How Reasons are Sensitive to Available Evidence', in C. McHugh, J. Way, and D. Whiting (eds.), *Normativity: Epistemic and Practical*. (Oxford: Oxford University Press), 90-115.
- Kiesewetter, B. (forthcoming), 'Are All Practical Reasons Based on Value?' *Oxford Studies in Metaethics* (Oxford: Oxford University Press).

²³ We want to thank Simon Kirchin and some anonymous referees for their comments. We also want to thank Maria Alvarez, Kyle Blumberg, Sam Carter, Christina Dietz, Claire Field, Simon Goldstein, John Hawthorne, Benjamin Kiesewetter, Maria Lasonen-Aarnio, Errol Lord, Martin Smith, Tim Williamson,

- Lazar, S. (forthcoming), 'Deontological Decision Theory and the Grounds of Subjective Permissibility', in M. Timmons (ed.) *Oxford Studies in Normative Ethics* (Oxford: Oxford University Press).
- Littlejohn, C. (2012), *Justification and the Truth-Connection* (Cambridge: Cambridge University Press).
- Littlejohn, C. (2020), 'Truth, Knowledge, and the Standard of Proof in Criminal Law', *Synthese* 197/12: 5253–86.
- Littlejohn, C, and Dutant, J. (2022), 'Even If It Might Not Be True, Evidence Cannot Be False', *Philosophical Studies* 179/3: 801–27.
- Lord, E. (2018), *The Importance of Being Rational* (Oxford: Oxford University Press).
- Mantel, S. (2017), 'Worldly Reasons: An Ontological Inquiry into Motivating Considerations and Normative Reasons', *Pacific Philosophical Quarterly* 98/1: 5–28.
- Moss, S. (2018), *Probabilistic Knowledge* (Oxford: Oxford University Press).
- Olsen, K. (2018), 'Subjective Rightness and Minimizing Expected Objective Wrongness', *Pacific Philosophical Quarterly* 99/3: 417–41.
- Papineau, D. (2021) 'The Disvalue of Knowledge', *Synthese* 198/6: 5311–32.
- Praolini, F. (2019), 'No Justificatory Closure without Truth', *Australasian Journal of Philosophy* 97/4: 715–26.
- Raz, J. (2010), 'Reason, Reasons and Normativity', in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 5 (Oxford: Oxford University Press), 5–23.
- Raz, J. (2009), 'Reasons: Practical and Adaptive', in D. Sobel and S. Wall (eds.), *Reasons for Action* (Cambridge: Cambridge University Press), 37–57.
- Ross, L. (2021), 'Rehabilitating Statistical Evidence', *Philosophy and Phenomenological Research* 102/1: 3–23.
- Ryan, S. (1991), 'The Preface Paradox', *Philosophical Studies* 64/3: 293–307.
- Sepielli, A. (2018), 'Subjective and Objective Reasons', in D. Star (ed.), *The Oxford Handbook of Reasons and Normativity* (Oxford: Oxford University Press), 784–99.
- Smith, M. (2018), 'When Does Evidence Suffice for Conviction?' *Mind* 127/508: 1193–218.
- Smith, M. (forthcoming), 'The Hardest Paradox for Closure', *Erkenntnis*.
- Steele, K.. (ms.) 'Crime, Punishment, and 'Individualized' Evidence'.
- Thomson, J. (1990), *The Realm of Rights* (Cambridge: MA.: Harvard University Press).
- Volokh, A. (1997), 'n Guilty Men', *University of Pennsylvania Law Review* 146/2: 173–216.
- Whiting, D. (2018), 'Right in Some Respects: Reasons as Evidence', *Philosophical Studies* 175/9: 1919–2208.
- Williamson, T. (2000), *Knowledge and its Limits* (Oxford: Oxford University Press).