n-1 Guilty Men[1]

Julien Dutant and Clayton Littlejohn

0. The Difficult Weekend

Agnes looks after n kids (i.e., 1, 2, 3, …, Inge, …, n-3, n-2, n-1, n). Over the course of a difficult weekend, she came to suspect each kid broke some rule. She suspected that n-9 stole something from 7, that 7 bullied n-14, and so on. When she suspects that one of the kids broke one of the rules, 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 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, Agnes would pause to think about the kid on the other side. 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. Build in what you must so that it seems to you, too, that if the relevant kid was guilty, Agnes knew it.

The plot thickens. Agnes's sister, Agatha, is known by all to be all but infallible. When she texted Agnes to say that precisely one of the kids was innocent, Agnes accepted that there were indeed 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, the follow up text never came. The

---

[1] Title inspired by Volokh (1997). For Volokh, 'n' is the number we solve for if we can either vindicate or improve upon Blackstone's suggestion that it would be better to let ten guilty men go free than to let one innocent person suffer the harms of punishment. (If Blacktone is right, n = 10.) For us, 'n' is the number of children (or men, women, etc.) being punished and our problem isn't that of finding the ratio of guilty to innocent people. (We think this question rests on a mistake.) Our question is whether we can punish n when it's known that n-1 are guilty.

[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.

network was down. So, it seems to us that she believes in n cases that a kid is guilty and knows, as we do, that she knows in n-1 cases that the relevant kid is guilty, while also knowing that one kid is innocent.

There are lots of questions we might ask about this case. 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, answers that we can use to test theories that purport to tell us which attitudes Agnes ought to have and which actions she ought to perform. We shall look at two ways of trying to vindicate our intuitions about these cases, a reasons-centred approach that's currently quite popular amongst the ethicists and an alternative that draws on decision-theoretic tools. We see something good in both approaches. Still, we think neither approach works quite right. We offer an alternative approach. It is a kind of hybrid. Our view has the virtue of being the only extant approach that makes sense of our hopefully not too idiosyncratic intuitions.

1. Should reasons be our guide?

According to one influential view, reasons should be our guide. This is the view of the *reasonologist*. Part of their idea is that normative reasons make the normative world go round.[3] If Agnes ought to X, she ought to X iff (and because) there is some reason (or reasons) that constitute a decisive reason for her to X.[4] Part of the idea is that normative reasons should provide guidance that the agent can follow. At a minimum, they have to be available to an agent so that she can succeed in responding as she should without having to make a stab in the dark or some deeply irrational choice.

Because reasons are taken by our reasonologists to be guides in both ways, this view would seem unpromising if their aim were to explain why Agnes *objectively* ought to do this or *objectively* shouldn't believe that. The philosophers who think that reasons should provide guidance that the agent can follow don't see themselves as offering an account of what we objectively ought to do, believe, or feel either because they are debaters who think that the genuine normative notions should be understood subjectively or prospectively or because they are dividers who

---

[3] Raz comes close to endorsing this when he says, "All normative phenomena are normative in as much as, and because, they provide reasons or are partly constituted by reasons" (2011: 85).

[4] This is the view developed and defended by Kiesewetter (2017) and Lord (2018). There are subtle differences between their views, but the differences won't matter here.

want to tell us what Agnes *prospectively* ought to do in a reasons-centred way.[5] Our concern is with the prospective reading of 'ought' unless stated otherwise.

To simplify our discussion, we assume that normative reasons are constituted by facts and that these are typically (but not exclusively) facts about the situation the agent faces (e.g., that 17 shoved 16, that n-20 put gum in Inge's hair, etc.).[6] 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 facts typically determine what duties we have, but that they might always have to be discernible to us:

> Even if duties are always grounded in features of the situation, it might be that to serve as a ground a feature 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 (2000: 57)?

We might imagine a spotlight that illuminates some facts and not others. The facts that can be seen by the agent that work together to determine how she should respond. The facts that aren't illuminated 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 take little interest in.

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.

---

[5] For a helpful overview of the difference between debaters and dividers, see Graham (2021) and Sepielli (2018).

[6] This is in keeping with the work of Alvarez (2010), Kiesewetter (2017), Lord (2018), and Mantel (2018). Some authors have argued that normative reasons might be falsehoods (e.g., (Comesaña, 2020)) or mental states, but we side with the above authors in rejecting the identification of reasons with mental states and argue against the use of false reasons in theories of justified response in Littlejohn & Dutant (forthcoming).

So far, so good. This fits with an ordinary way of talking about the difficulty duties of being a parent of difficult children. 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 that helps.)

This explanation covers a very small part of *The Difficult Weekend*. How do we cover the rest of it? 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 [name] to their room to think about what they've done because [fact].

For each kid that Agnes ought to believe was guilty of breaking a rule, we'd have to fill this in:

> Agnes should believe [name] broke a rule. This is because she has a decisive reason to believe this: it's the fact that [fact].

For each kid that Agnes ought to punish, we'd have to fill this in:

> Agnes should punish [name]. She has a decisive reason to do so: it's the fact that [fact].

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].[7]

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.[8] 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

---

[7] While we think that Known Guilt is *plausible*, we think that it will need to be qualified and interpreted properly. We shall discuss Known Guilt further in §3.

[8] A more demanding epistemic relation could be substituted for knowledge, but we should remember that this relation has to be one that we plausibly bear to facts about the situation. It's hard to imagine that some such relation would avoid the issues discussed here.

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 have to be informed by independently plausible claims about what Agnes can know, fit with what we know about the prospective 'ought', and their description has to 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. We think the arguments cannot each be sound because their conclusions are jointly inconsistent. We don't think the reasonologists can reject the assumptions that generate this inconsistency. The lesson that we take from this is that the reasonologists are using the wrong tools to handle *The Difficult Weekend*.

Let's consider our first argument:

> *The Not-All Argument*
>
> 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 a tune, to doodle, or to do a handstand in order for it to be the case that it's fine for her to do such things. Everyone who thinks normative reasons explain oughts has to agree that when 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 is might be like that. Punishment is another. Being disappointed in someone for what they've done is yet another. In each instance, the absence of a good reason to respond 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 will ensure the presence of sufficient reasons not to X.) While carrying an umbrella is probably not the sort of thing that's 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 (P2na) 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*
> P1sm. 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.
> P1sm. 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.).
> Csm. So, some kids must be punished and must be blamed.

Remember that our reasonologist takes knowing to be sufficient for possessing a reason and the facts known to Agnes seem like reasonable candidates for being sufficient or decisive reasons to blame and punish. If Agnes were to refrain from blaming and punishing when she knew that the relevant kids were guilty of serious wrongdoing, she would systematically fail to respond to these possessed reasons.

Here's our third argument:

> *The All-Or-None Argument*
> P1aon. 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.
> P1aon. But there is no such difference between the cases.
> Caon. 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.[9] Our reasonologists accept this

---

[9] A classic case for this is found in Cohen (1984). Our reasonologists accept this constraint.

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 just stipulate. Agnes reached her n conclusions about guilt by relying on the right kind of evidence in n cases. She processed that 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 reasonologist say in response?

First Response: Some Children Must be Punished
The reasonologist 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*. 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.

Second Response: If Many Were Guilty, All Should be Punished
Let's consider an alternative response. The reasonologists might reject the conclusion of the *The Not-All Argument*. 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 reasonologists pursue 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 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].

Kiesewetter (2017) and Lord (2018) have proposed that back-up reasons (i.e., facts known to the agent in the bad case that differ from the ones the agent would cite in the good case in defending her responses) 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.

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.

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:

- That Inge bullied n-8;
- That Inge probably did something wrong/probably bullied n-8;
- Most of the kids did something bad that warrants punishment, blame, etc.;
- 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.

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. But 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?) As a substantive point, 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* 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.).[10] 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). Shakespeare didn't choose to tell his tale with such a character. Othello's reasons had to do with fidelity, not the appearance of it, because he valued fidelity, not the appearances of it.[11] We should not generally assume that if the fact that *p* constitutes a decisive reason for an agent who knows *p* to X, 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 X. As the case of emotion

---

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

[11] 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.

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?[12]

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.

It seems that 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.

1.4 Summing Up

We do not believe that there is any set of facts (if it includes facts about the situation) that could directly determine how Agnes ought to respond. The reasonologist 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

---

[12] This is the line that Smith (2018) takes. He thinks that in preface cases, we cannot have evidence that provides sufficient support for each proposition. 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 each. 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. See Littlejohn & Dutant (2021) for discussion of the preface and Littlejohn (2020) for discussion of punishment and blame in such cases.

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 a prospectivist says Agnes ought to punish or ought to free a kid.

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.[13] According to one version expectabilism, A ought to X iff (and because) X-ing maximises expected value. 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 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 these values and the probabilities let us rank Agnes's options in terms of expected desirability. 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 prospectivists like.

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
Refrain	c	d

Agnes thinks about the possible outcomes (a-d), knows their values, and is able to 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

---

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

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 in spite of the fact that it's highly probable in each case that the kid is guilty.[14] 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 view we've formulated.

We see a number of problems for the expectabilist view as it's been formulated. In addition to the fact that we don't think naked statistical evidence warrants blame, belief, or punishment, we don't see how the expectabilist can explain the connections between these responses. We think that it would be wrong to punish when it's wrong to blame or when it's wrong to believe someone to be guilty. This is because punishment should be, *inter alia*, a way of expressing blame and holding responsible. In general, Agnes is in no position to express blame or hold someone responsible for her bad behaviour if she knows she doesn't know that the relevant party has done anything wrong. The expectabilist view doesn't support this requirement. In cases like *The Garden Party*, Agnes can be certain that she doesn't know if Inge has done anything wrong, but she nevertheless is told she should punish Inge in spite of her ignorance. We think that it's absurd to say that someone should blame someone when they know they don't know the relevant party has done anything wrong and unreasonable to say that we should punish when we should refrain from blaming and refrain from holding responsible.

3. What could our guide be?

Reasonologists cannot vindicate intuitions about the ways that Agnes prospectively ought to respond in *The Difficult Weekend*. Here are two diagnoses. They need

---

[14] We side with Adler (2002), Buchak (2014), Enoch et al. (2012), Littlejohn (2020), Moss (2018), Smith (2018), and Thomson (1986) 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.

probabilities to deal adequately with preface cases. [15] Once they're introduced, however, we'll explain what Agnes ought to do in terms of the probabilities that various reason-constituting facts obtain (and not in terms of sets of such facts, possessed or otherwise). The expectabilist does a better job with *The Difficult Weekend* but struggle with *The Garden Party*. Here's our diagnosis. They can only get the cases right if they work with the right tools for characterising what's objectively desirable. It's here that they need reasons and not (just) values. Thus, one lesson that we'd take from this is that there are some (but only some!) cases in which we'll need the tools from both approaches. The decision theorists will sometimes need reasons and the reasonologists need more decision theory.

Here's how we explain the verdicts. Let's start with this:

Known Guilt: Agnes ought to believe, blame, and punish
if she's in a position to know that the relevant kid is guilty.

When it comes to believing, blaming, and punishing, it matters that the subject can identify good reasons to blame, to believe, and to punish. There is *nothing* desirable about blaming someone in the certain absence of available reasons to blame. Knowledge of wrongdoing should suffice for the possession of such reasons.

We also think that knowledge of wrongdoing is necessary for the possession of the reasons that can justify believing, blaming, and punishing. 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:

Only Known Guilt: Agnes should not believe, blame, or
punish if she doesn't know the relevant kid is guilty.

When you know you don't know someone has done something wrong, you cannot hunt around for other reasons to blame, facts that might obtain when the target of blame has, so far as you can see, done nothing wrong.[16]

Won't combing Known Guilt and Only Known Guilt lead us back to the problem that Agnes ought to punish some but only some of the kids? Yes. But this is only a problem if we understand these claims as claims about *prospective* rightness. We take them to be claims about *objective* rightness. Knowledge of guilt is necessary and sufficient in our cases for objectively proper blame, belief, and punishment. It's here that the reasonologist's reasons matter, in explanations of objective rightness.[17]

---

[15] But not just to deal with preface cases. They also need them to make sense of relatively uncontroversial observations about sweetening and souring by increasing or decreasing the chances of success. See Littlejohn (2019).

[16] Note that the cases involving blame and punishment are exceptional because there are crucial normative reasons. It's not generally true that it's only right to perform an act if a certain kind of reason exists and so it's not surprising that it's not generally true that it matters whether certain facts are known to the agent.

[17] If this is right, the idea that reasons must be a guide needs attention. Reasons can still be a guide in establishing guidelines understood as lines that we cannot

Preface cases show that sets of known facts aren't the right tool for telling us what agents *prospectively* ought to do. Here, we need the expectabilist's probabilities and expectations. However, given that reasons matter to objective rightness, it's not right that reasons don't play any role in the expectabilist story. If we know that some factors matter to objective suitability, the likelihood of their presence and absence should matter to prospective rightness. Given Known Guilt, it would be wrong (and so undesirable in one sense) to fail to punish and blame those known to be guilty. Given Unknown Guilt, it would be wrong (and so undesirable in that sense) to punish those not known to be guilty. Prospectively speaking, we should aim to minimise expected undesirability.[18] Reasons tell us what's desirable and undesirable (objectively speaking). Probabilities then tell us, prospectively speaking, how we should strive to conform to reasons given our doubts about their presence and absence.

If we're right and there are normative reasons to conform to Known Guilt and Only Known Guilt that apply to us even when these reasons aren't known to us, we can offer a view that's similar in its mechanics to the expectabilist views we're familiar with already. Instead of focusing on the difference in *value* between possible outcomes, we focus on the difference in the *weight* of reasons for or against performing acts of a type or bringing about some outcomes. If the weights of reasons can be quantified in the ways that values must be for the expectabilist view to work, we can use the tools of decision theory and the reasonologists for handling our two tricky cases. Let's say that it's wrong to some degree to violate Known Guilt, wrong to some degree to violate Only Known Guilt, and it is wrong to a greater degree to violate the latter than the former. This would give us this proposal:

> Objectively speaking, Agnes ought to punish, believe, and
> blame iff she's in a position to know the relevant kid is guilty.
> Prospectively speaking, Agnes ought to punish, believe, and
> blame only if it's more likely than not that she's in a position
> to know the relevant kid is guilty.

When it comes to *The Difficult Weekend*, we think that (assuming n is suitably large) it's sufficiently likely in each case that Agnes knows the relevant kid to be guilty. This holds in each of the n-cases, so we get that if Agnes ought to punish any, she ought to punish each. She's keenly aware that she's at risk in each case of blaming a kid without being able to offer a reason that justifies it. Thus, she's wrong that there's

---

permissibly cross. They won't provide guidance beyond this, but the preface case shows that neither possessed nor unpossessed reasons were fit for this role. If we can be rationally uncertain about which reasons obtain (which we can if we can sometimes identify them with facts about the situation), guidance comes from hypotheses about which reasons obtain and facts about their weights if they should obtain.

[18] For recent exploration and defence of these ideas, see Lazar (2020) and Olsen (2018).

a risk of objective wrongdoing. She's also keenly aware that there's a risk of failing to blame the blameworthy. Given the magnitudes, it seems she should take the risk once, twice, thrice, etc. When it comes to *The Garden Party*, there is *no* chance that she's aware of a suitable reason to blame or punish. She's aware that such reasons are 'out there', as it were. We think it's objectively wrong for her to blame first and find the reason later. Since the probability of conforming to Only Known Guilt is effectively zero, she should refrain.

4. Conclusion

We want to conclude this discussion with a brief discussion of what we hope future work on normative reasons will be. We think that reasonologists should use decision-theoretic resources and that decision-theorists should take an interest in reasons.

Our approach to our cases involved three elements. First, it includes a substantive element, the claim that knowledge matters to (objectively) suitable belief, blame, and punishment. We focused on punishment because, we think, proper punishment expresses attitudes and so differs from humdrum cases of deciding whether to take an umbrella or mix this stuff with that stuff to make drinks for the guests. This gives you, if you like, our substantive story about when responses of certain types (i.e., blame and punishment) are suitable and how their suitability is bound up with the suitability of further attitudes (i.e., beliefs about wrongdoing). It's here that we think facts that constitute normative reasons can earn their keep. It's *not* part of our view that knowledge or possession matters whenever we're trying to describe what's objectively suitable. If we were talking about saving miners or investing in stocks, it's not part of our view that, say, whether it's objectively right to try to save the miners this way or that way depends upon what the subject is in a position to know. For us, knowledge matters in our cases because punishment is supposed to express attitudes that would be suitable only if the agent is cognisant of facts about wrongdoing. It should not be taken as a general proposal about the importance of knowledge to right action.

Second, we found some role for reasons to shape our understanding of objective desirability or suitability. We doubt that we can characterise what's objectively desirable about having just the beliefs we objectively should or blaming and punishing the people we should in purely evaluative terms.[19] We also doubt, however, that when we shift the from evaluative conceptions of desirability to some reasons-centred story about desirability, the decision theorists should be spooked. Reasons and their weight can function formally like values, so the move from values

---

[19] See Côté-Bouchard (2017) and Raz (2011) for arguments that some normative reasons aren't value-based.

to the weights of reasons feels somewhat notational.[20] We think the standard expectationsit views (e.g., those that tell us that we ought to punish if it's sufficiently likely that someone is guilty) are mistaken because they're concerned with expectations of guilt, not because they're concerned with expectations. We think they should focus on expectations of wrongfulness and that the problems emerge for the view because everyone seems to assume that we should be concerned with expectations iff we should think the issue is settled by expectations of guilt. We reject that, but that gives us freedom to explore the idea that other expectations matter to prospective rightness.

Third, we think that the theory of prospective rightness needs to draw from some substantive claims about what objective rightness consists in and then use the tools of decision theory to tell us how to cope with uncertainty about right-making features. This, we think, is what the reasonologists have missed. They were right that something about the agent's epistemic state matters to prospective rightness, but they were wrong to think that *binary* relations (e.g., belief, knowledge) between existing facts and the agent was the thing to focus on. We need something that comes in *degrees* and to recognise that once we shift to a degreed notion, what matters is not a degreed relation to a fact, but degreed relations to possible facts.[21]

We're not quite sure why this is, but reasonologists seem to be generally uninterested in using the tools of decision theory to explain what we prospectively ought to do. We don't see how they can, without using these tools, handle *The Difficult Weekend* or explain why the size of n should matter to our intuitions in variations of this case. We think that when n gets sufficiently small, there might be cases where it's possible for Agnes to know that some kids are guilty where, intuitively, it wouldn't seem reasonable for her to blame and to punish. Again, this would be difficult, if not impossible, to explain if we approached this case using sets of possessed reasons/known facts about guilt, but easy to explain using expectationist tools. Even if it's worse to violate Only Known Guilt than to violate Known Guilt, if the risk of violating Only Known Guilt is sufficiently high, we can

---

[20] One important difference, however, is that reasons for or against believing don't seem to give us reasons to accept trade-offs. This speaks to the possibility of consequentialising a theory of suitable belief and blame.

[21] Our proposal differs from Whiting (2018) who identifies reasons with evidence of right-makers. We identify objective normative reasons with the right-makers and see no need to introduce further kinds of normative reasons into our theory of prospective rightness. Thus, we can avoid the objection, often directed against Kearns and Star's (2009) 'reasons as evidence' view, that too many things would be identified as normative reasons if anything that boosted the probability of certain normative propositions counted as a reason. Since we've tried to show that we don't need sets of propositionally specified reasons to determine what the agent prospectively ought to do, we also think we've undermined a key motivation for introducing false normative reasons.

envisage cases (e.g., where the number of known innocent kids increases) where we'll judge both that Agnes violates Known Guilt if she doesn't punish *and* judge that that's what she should do.[22] We think that future work on reasons should explore ways to use the tools of decision theory or find reasons for resistance.

Some authors have suggested that there are reasons to resist using the tools of decision theory in cases like ours. Buchak (2014) thinks that in cases like *The Garden Party*, we shouldn't punish and that we should reject the idea that the high probability of guilt should be sufficient for blame and punishment. She sees in this, though, problems for expectationism:

> … the natural home of credence is in consequentialist norms, and the natural home of belief—and the domain in which we cannot eliminate belief in favor of credence— is in deontological norms (2014: 306).

We are more optimistic. While we agree that the probability of guilt is *not* the thing that explains why Agnes should or shouldn't punish or blame, we think that a view on which rational credences about what we should believe or can know tracks our intuitions about this case. When it's nearly certain that Agnes knows, we think people generally agree she prospectively should punish. When it's quite likely that she couldn't know, we think people generally agree she prospectively shouldn't punish. If, as we conjecture, there are possible cases in which she knows some specific kids to be guilty but the probability that she knows isn't sufficiently high, we think it might seem right that Agnes prospectively shouldn't punish. We don't see how to understand scale effects (e.g., the importance of the known ratio of success to failure) and the importance of belief to blame without taking account of *both* rational credence and belief (i.e., credences about the conditions that determine when belief and blame would be objectively suitable).

In summary, we hope that this exercise has shown one of the ways in which we can solve difficult puzzles by combining the ideas of reasonologists and decision-theorists. We hope future work continues to explore the strengths and limits of this kind of hybrid approach.

References

Adler, J. E. (2002). *Belief's Own Ethics*. MIT University Press.

Alvarez, M. (2010). *Kinds of Reasons: An Essay in the Philosophy of Action*. Oxford University Press.

Buchak, L. (2014). Belief, Credence, and Norms. *Philosophical Studies*, *169*(2), 1–27.

Christensen, D. P. (2004). *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford, England: Oxford University Press.

Cohen, S. (1984). Justification and Truth. *Philosophical Studies*, *46*(3), 279–295.

---

[22] Our argument assumes the possibility of improbable knowledge. See Williamson (2014) for a defence of this possibility. See Lasonen-Aarnio (2010) and Littlejohn & Dutant (2021) for a discussion of the possibility of unreasonable knowledge.

Comesaña, J. (2020). A Plea for Falsehoods. *Philosophy and Phenomenological Research*, *100*(2), 247–276. https://doi.org/10.1111/phpr.12523

Côté-Bouchard, C. (2017). Is Epistemic Normativity Value-Based? *Dialogue, 56*(3), 407–430.

Dancy, J. (2000). *Practical Reality* (Vol. 78, Issue 305, pp. 414–425). Oxford University Press.

Enoch, D., Spectre, L., & Fisher, T. (2012). Statistical Evidence, Sensitivity, and the Legal Value of Knowledge. *Philosophy and Public Affairs*, *40*(3), 197–224.

Graham, P. A. (2021). *Subjective Versus Objective Moral Wrongness*. Cambridge University Press.

Kearns, S., & Star, D. (2009). Reasons as Evidence. *Oxford Studies in Metaethics, 4*, 215–242.

Kiesewetter, B. (2017). *The Normativity of Rationality*. Oxford University Press.

Lasonen-Aarnio, M. (2010). Unreasonable Knowledge. *Philosophical Perspectives*, *24*(1), 1–21.

Lazar, S. (2020). Duty and Doubt. *Journal of Practical Ethics*, *8*(1), 28–55.

Littlejohn, C. (2019). Being More Realistic About Reasons: On Rationality and Reasons Perspectivism. *Philosophy and Phenomenological Research*, *99*(3), 605–627.

Littlejohn, C. (2020). Truth, knowledge, and the standard of proof in criminal law. *Synthese*, *197*(12), 5253–5286.

Littlejohn, C., & Dutant, J. (2021). Defeaters as Indicators of Ignorance. In M. Simion & J. Brown (Eds.), *Reasons, Justification, and Defeat*. Oxford University Press.

Littlejohn, C., & Dutant, J. (forthcoming). Even if it might not be true, evidence cannot be false. *Philosophical Studies*, 1–27.

Lord, E. (2018). *The Importance of Being Rational*. Oxford University Press.

Mantel, S. (2018). *Determined by Reasons: A Competence Account of Acting for a Normative Reason*. Routledge.

Moss, S. (2018). *Probabilistic Knowledge*. Oxford University Press.

Olsen, K. (2018). Subjective Rightness and Minimizing Expected Objective Wrongness. *Pacific Philosophical Quarterly*, *99*(3), 417–441.

Papineau, D. (2021). The disvalue of knowledge. *Synthese*, *198*(6), 5311–5332.

Raz, J. (2011). *From Normativity to Responsibility*. Oxford University Press.

Ross, L. (2021). Rehabilitating Statistical Evidence. *Philosophy and Phenomenological Research*, *102*(1), 3–23.

Sepielli, A. (2018). Subjective and Objective Reasons. In D. Star (Ed.), *The Oxford Handbook of Reasons and Normativity*. Oxford University Press.

Smith, M. (2018). When Does Evidence Suffice for Conviction? *Mind*, *127*(508), 1193–1218.

Steele, K. (n.d.). *Crime, Punishment, and 'Individualized' Evidence*.

Thomson, J. J. (1986). Liability and Individualized Evidence. *Law and Contemporary Problems*, *49*(3), 199–219.

Volokh, A. (1997). *N* Guilty Men. *University of Pennsylvania Law Review, 146*(1), 173.

Whiting, D. (2018). Right in some respects: Reasons as evidence. *Philosophical Studies, 175*(9), 2191–2208.

Williamson, T. (2014). Very Improbable Knowing. *Erkenntnis, 79*(5), 971–999.