

On the Normativity of Rationality and of Normative Reasons

Julien Dutant (King's College London)

Clayton Littlejohn (King's College London)

Abstract: Scepticism about the normativity of rationality is often partially based on the assumption that normative reasons are normative. Starting from the assumption that normative reasons are normative, someone will argue that reasons and rationality can require different things from us and conclude that rationality must not be normative. We think that the assumption that normative reasons are normative is one that deserves more scrutiny, particularly if it turns out, as we shall argue, that no one has yet shown that the requirements of reasons and rationality might be unified. We look at the most promising proposals about how to unite the requirements of reasons and rationality and argue that they cannot succeed. In the course of doing so, we discuss a case that we think gives us good reason to think that good reasons are not the things that determine in each case what we ought to believe and/or do. We argue that the best way to deal with the examples discussed here might be to appeal to principles of rationality and acknowledge that following their guidance will not invariably ensure that we do what normative reasons (as they are usually understood) support. We think that it might just follow that normative reasons (as they are usually understood) cannot determine what would be rational to do or what we ought to do.

1. Introduction

There is now considerable scepticism about the normative significance of rationality:

Normativity of Rationality: A ought to ϕ iff A is rationally required to ϕ .¹

Some of this scepticism seems to derive from the idea that normative reasons and rationality can require different things from an agent.² The thought that rationality might require us to ϕ when there is no reason to ϕ seems to have convinced many that this thesis is mistaken:

Unification: A is rationally required to ϕ iff there is a reason for A to ϕ that is a decisive reason for A to ϕ .³

¹ Some authors distinguish substantive rationality from some non-substantive notion (e.g., procedural rationality). If substantive rationality is understood in a certain way and is taken to be the operative notion of rationality, this thesis might seem trivial. If, however, it is concerned with rational requirements that do not appear substantive, the thesis is much more interesting and much more controversial. See Kiesewetter (2017) and Lord (2018) for a careful attempt to ground the requirements of rationality that do not appear on first inspection to be substantive in terms of the demands of reasons. It is this kind of project that interests us here.

² See, for example, Broome's (2005, 2007) argument against the idea that rationality might consist of responding correctly to reasons. We appreciate that there are other grounds for scepticism about the normativity of rationality. In this discussion, we want to focus on a particular kind of challenge to the Unification Thesis.

³ While some authors (e.g., Schroeder (2008) and Whiting (2014)) have tried to characterise rationality in terms of subjective reasons, the Unification Thesis is concerned with the reasons that determine what the agent ought to do in Broome's (2013: 26) 'central' sense. Fans of subjective reasons accounts of rationality might be pleased to see that we agree that we cannot understand rationality in terms of responding correctly to reasons that are distinct from subjective reasons (i.e., because some of them are identified with objective matters of fact). They will not be pleased,

To the extent that scepticism about the normativity of rationality is grounded in concerns about the Unification Thesis, this scepticism is partially based on the assumption that normative reasons are normative:

Normativity of Reasons: A ought to ϕ iff there is a reason for
A to ϕ that is a decisive reason for A to ϕ .⁴

An argument that, say, an agent might be rationally required to ϕ when there is not a single good reason to ϕ only threatens the idea that rationality is normative if we assume that every genuinely normative requirement is generated by normative reasons.

We think that this assumption about normative reasons should be made explicit and critically examined. While we don't think that the claim that normative reasons are normative is *prima facie* implausible, we also don't think that the claim that rationality is normative is *prima facie* implausible. As Lord (2018) recently observed in a discussion of the normativity of rationality, the claim that we ought to be rational seems platitudinous. If both the normativity of rationality and the normativity of reasons seem rather plausible, we don't think it's obvious that the right response to a convincing objection to the Unification Thesis is a rejection of the normativity of rationality. If we cannot find a way to square the normativity of reasons with the normativity of rationality, perhaps it's time to abandon this idea that normative reasons determine what we ought to do.

We don't deny that the normativity of reasons is a *prima facie* plausible operating assumption, but we don't think that it should be treated as a fixed point that cannot be revised. The question, 'Are normative reasons normative?' seems like a closed question, but it is a sensible one to ask when the context makes it clear that we are operating with some specific and controversial view of reasons. People now tend to think of normative reasons as facts, typically the facts about the situation. When the context makes it clear that the question, 'Are normative reasons normative?' is really just the question, 'Do features of the agent's situation determine what the agent should do?' the question is just Prichard's (2002) question. We think this is a good question to ask.

however, to see that our criticism of the Unification Thesis causes trouble for subjective reasons accounts of rationality.

⁴ Readers need to take some care in reading this thesis. Some authors will say that there can be things that are among the reasons there are that do not have any bearing on what the agent ought to do. They will say that only the possessed reasons or the reasons that the subject has determine what this agent ought to do. We think that it might be helpful to think of things in terms of *potent* reasons, the reasons that do help determine what the agent ought to do. Let us assume that this restriction is in place. Readers who think that only possessed reasons are potent need not take issue with this thesis. We will have more to say about possession below.

To us, it seems that many people working on reasons are moving away from the view that large collections of facts that the agent couldn't have known about determine what an agent ought to do towards a view on which facts about an agent's perspective help to determine what she ought to do.⁵ If nothing else, we think that it's clear that anyone who accepts a pure objectivist view of 'ought' will reject the Unification Thesis and deny that rationality is normative. We think that it makes sense, then, to focus our discussion on a view we'll call *reasons perspectivism* because we think that it represents the most promising option for someone who wants to reconcile the normative of normative reasons (understood as facts that would typically be facts about the agent's situation) with the normativity of rationality. If we can show that this view doesn't give us the resources we need to defend the Unification Thesis *and* that this view doesn't deliver verdicts that fit our intuitive judgments about cases, we think that it might be right to conclude that normative reasons (on their usual understanding) are not normative. In §2, we shall flesh out the reasons perspectivists' view and explain how they have tried to defend the Unification Thesis. In §3, we shall argue that the reasons perspectivists' normative reasons are not normative. In §4, we shall critically discuss some strategies for trying to deal with our cases. We survey what we take to be the most promising alternatives to reasons perspectivism and argue that each of these views delivers the wrong verdicts about our cases. In §5, we shall propose two ways of trying to handle our case. While we think that these proposals have important virtues that rival view lacks, both views predict that the requirements of rationality can come into conflict with the demands of normative reasons.

2. Reasons, Rationality, and the Agent's Perspective

To ensure that we're all up to speed on the recent debates about the Unification Thesis, we shall quickly review the shared operative assumptions about reasons and rationality and explain how the reasons perspectivists have responded to Broome's argument against the thesis.

Let's start with reasons. On a now standard view, normative reasons are the kinds of things that count in favour or against, something we would think of as a pro or a con, or something that we would weigh up in deliberation when considering countervailing considerations that seem to support alternative options.⁶ We quite like the metaphor of the ledger, a thing where you would list the pros and the cons and either count them (as Benjamin Franklin would) or weigh them

⁵ Objectivism still has its defenders, however. See Graham (2010), Littlejohn (2012), and Thomson (1990) for defences of objectivist views about 'ought' and/or obligation.

⁶ Way (2017) takes reasons to be premises of good reasoning. Gregory (2016) proposes that they are good bases for action. Mantel (2017) characterises them as the states of affairs that premises of practical reasoning concern. This is all in keeping with the idea that normative reasons are facts, states of affairs, or true propositions. See Alvarez (2010) and Dancy (2000) for defences of the idea that reasons of all kinds are facts.

against each other (as most reasonable people would).⁷ If we think of reasons in this way, it's natural to think of them as facts or true propositions. They would typically be facts about the situation (e.g., the fact that you promised, the fact that it hurts to be kicked, the fact that they need help, etc.). Let's call this view *factualism*.

Factualism is a claim about what reasons *are*, not what they *do*. What do reasons do? The factualists who believe normative reasons are normative think that facts can constitute normative reasons when they count in favour of doing something. These reasons/facts have weights that we should consider in deliberation. They help make a case for something. Without worrying too much about the different ways of understanding how reasons support things, it's important to remember that if the Normativity of Reasons Thesis is correct it must be that whenever some agent ought to ϕ there is some fact or some facts that constitute at least one reason to ϕ . If an agent ought to ϕ but there is not at least one fact-constituted reason to ϕ , we would either need to accept some alternative to factualism or we would have to concede that the Normativity of Reasons Thesis is false. A sensible strategy for attacking the idea that normative reasons are normative is find cases where an agent ought to ϕ even though there is not one single reason for this agent to ϕ .

We face an obvious challenge in trying to find such cases. Whenever we consider a case where it's clear that some agent ought to ϕ , we know that it seems plausible to say that there must have been some reason to ϕ and to say that there must be some fact or facts that explain why this agent ought to ϕ . Doesn't the first point suggest that there must exist at least one reason for an agent to ϕ when this agent ought to ϕ ? Doesn't the second suggest that such a reason must be some fact or facts? Doesn't that support factualism?

We think that it's natural to conclude that there is 'some reason' for an agent to ϕ whenever we have decided that this agent ought to ϕ . We also think that it's plausible that there must be some fact or facts that explain why this agent ought to ϕ if indeed this agent ought to ϕ . What we have to guard against, however, are two potential slips. First, we would like to flag Fogal's (2016: 81) note of caution.⁸ Much of our 'reason' talk might involve using 'reason' as a mass noun instead

⁷ See Fantl and McGrath (2009: 77). The ledger metaphor might be too simplistic in some ways. Thinking of things in terms of listed pros and cons might not be the best way to capture the interconnections between reasons that in isolation would seem to favour but appear to count against when taken in combination (e.g., the fact that a dish contains chocolate might count in favour of eating it and the fact that it contains mashed potatoes might count in favour of it, but they do not combine to favour eating). Still, we think that it helps us see what it would be for a reason to enter into a kind of practical deliberation.

⁸ For further critical discussion of the idea that each 'ought' involves one or more normative reasons, see Littlejohn (forthcoming) and Titelbaum (forthcoming).

of a count noun. It isn't clear that whenever there is sufficient reason (mass) to ϕ , there must be one or several reasons (count) to ϕ . Second, we would remind readers that there might be no interesting relationship between (a) things that figure in explanations of normative facts and (b) things that constitute pros or cons, facts that we weigh up in deliberating about whether to ϕ , things that would be listed in our ledgers, or things that would constitute a good basis for deciding, choosing, and so on.⁹ As we shall see below, some of the philosophers who defend factualism and defend the idea that normative reasons are normative think that the things that determine what an agent ought to do are (a) facts that count in favour in familiar ways (i.e., by being pros) and (b) facts that the agent knows. It isn't clear that in every possible case where some agent ought to ϕ there is some fact or some facts that meet both of these conditions. Indeed, we think that many facts that figure in familiar normative explanations will *not* meet both conditions.

When it comes to rationality, all parties to this discussion accept *mentalism*, the view that rationality supervenes upon an individual's non-factive mental states.¹⁰ This doesn't tell us much about rationality because mentalism is silent on further questions about rationality that matter later. Let's mention three. The first concerns knowledge and rational belief. No non-sceptical mentalist thinks that it's rational to believe only what you know, but many would agree that anything we know is something that we rationally believe. Knowledge is sufficient for rationality. Second, when it comes to evidence, many would agree with Kieseewetter (2018) that a thinker is rationally required to believe p when she is actively considering whether p and she possesses sufficient evidence for this belief. Third, it seems that there are rational connections between some beliefs and some intentions and actions. If a thinker believes that she ought to ϕ , it would seem that it would be irrational for her to intend to do other than ϕ or to intentionally do other than ϕ . If the akratic agent is irrational, it would seem that there is a requirement on rational agents to see to it that they do not both believe that they ought to ϕ and do otherwise. This is the Enkratic Requirement.

These assumptions tell us how we should understand claims about the normativity of rationality and of reasons. Let's now consider Broome's (2005) objection to the Unification Thesis.

⁹ Kieseewetter (forthcoming) criticises Wedgwood's (2017) recent work on rationality on the grounds that Wedgwood's account severs the connection between rationality and reasons. Kieseewetter argues from the point (which we accept) that whenever there is some ought there is some fact or facts that explain it to the conclusion (which we are sceptical of) that whenever there is some ought there is at least one fact that constitutes a reason to do that thing. We think that the fact in question might not be the kind of fact that does what people say reasons have to do to be reasons (e.g., count in favour, constitute a good basis, figure as a premise in reasoning, carry some weight, etc.).

¹⁰ For a defence, see Cohen (1984) and Wedgwood (2002).

The salmon smells delicious. It is high in protein and low in cholesterol. It is rude not to eat what you've been served. These facts all count in favour of eating. And since this is all you know about the fish, it seems that it's rational to eat it. What you don't realise is that the fish contains salmonella. Broome thinks that this is a very good reason not to eat it. It outweighs all the reasons to eat it. So, it seems that what you ought to do is throw the fish away even though it's not rational to do this. So, we are supposed to conclude, the Unification Thesis is mistaken.

Kiesewetter (2017) and Lord (2018) have pointed out that this objection to the Unification Thesis only works if we assume that reasons the agent could have known nothing about help determine what the agent ought to do. But this objectivist assumption about the facts/reasons that determine what we ought to do is not one that the factualist needs to accept. If we thought that the fact that the salmon contained salmonella constituted a decisive potent reason for the relevant agent on that occasion not to eat, then it would make it the case that the agent ought to throw out the fish. The clash between reasons and rationality would be unavoidable. If, however, we think that potent normative reasons (i.e., those that help to determine what a particular agent in some particular situation ought to do) are epistemically constrained in such a way that only the accessible ones are potent, Broome's objection fails. The accessible reasons in his example all count in favour of eating, so they all count in favour of doing the thing we all agree is rational. It's clear that Broome's objection misses its mark once we abandon the objectivist assumptions about 'ought'.

Dancy floated the idea of imposing an epistemic restriction on potent reasons:

...features [of the situation] 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).¹¹

¹¹ We do not think that Dancy endorses the idea that the potent reasons are the reasons that happen to make it through the epistemic filter (i.e., the facts that the agent knows or is in a position to know). We think this in part because he said as much in personal communication. He also does not want to defend the Unification Thesis. He thought that instead of thinking of the potent reasons as the evidence-relative reasons (i.e., the things we know or are in a position to know), we should think of them as what we know or *should* have known. It wasn't clear whether this suggestion was that the things not known that should have been known were, as Goldberg (2017) conceives of it, an epistemic failure (albeit connected to some social practice) or some sort of non-epistemic failure. We want to note two things. First, it is unlikely that someone who works with a normative conception of the things that constitute normative reasons will defend the central theses discussed in this paper (e.g., that all the oughts are ultimately grounded in demands generated by potent reasons). Second, we think that the example discussed below should be a problem for views that identify the potent reasons with the facts that the agent knew or should have known because

Reasons perspectivists take this idea and run with it. They agree that what we ought to do on any particular occasion is determined by reasons understood as facts (and that these will typically be facts about the situation). They also agree that what we ought to do and what we would be rational to do are the same thing. To ensure that reasons and rationality do not come into conflict, they impose epistemic constraints on potent normative reasons to ensure that all the reasons that help to determine whether some individual ought to ϕ are accessible to that agent:

Evidence-Relativity: The reasons that determine whether A ought to ϕ are the facts that are part of A's evidence (i.e., those that A knows or is in a position to know).¹²

By limiting the potent normative reasons to the possessed or accessible reasons, reasons perspectivists hope to undercut Broome's objection to the Unification Thesis. As reasons perspectivists see things, whenever rationality requires something from an agent, this *is* what the agent ought to do and it is what the agent ought to do because this is something that the evidence-relative reasons require. The idea that epistemic rationality is just a matter of responding correctly to the thinker's evidence is a popular one. We can think of reasons perspectivism as a generalisation of this idea.

Before we turn to our critical examination of their view, let's briefly consider the main motivations for the reasons perspectivists' idea that the facts that determine what an agent ought to do are those that constitute evidence-relative reasons. We see three main motivations for this view. The first simply appeals to the success of reasons perspectivism. Many believe that it's plausible that the factualist's normative reasons are normative and think that it's plausible that we ought to be rational. If a view can vindicate both of these ideas, there is much to be said for the view.

The second appeals to an influential critique of pure objectivist views. On pure objectivist views, there are no epistemic constraints on the factors that determine whether A ought to ϕ . Unknown facts and unforeseeable consequences might be in the mix (e.g., the facts about salmonella or the terrible consequences of eating contaminated fish might make it the case that the agent ought not eat Broome's delicious salmon). Many believe that our intuitions about three-options cases show that these pure objectivist views must be mistaken.¹³

we don't think that our protagonists should have known that the innocent child was innocent. Readers will have to decide whether they agree with us on this point.

¹² See Kiesewetter (2017, 2018) and Lord (2018) for defence of roughly this proposal.

¹³ Regan (1980) presents first three-option case. The case has figured prominently in criticisms of objectivist views of ought and obligation. See Jackson (1991), Kiesewetter (2017), Lord (2018), and Zimmerman (2008).

Here is a familiar three-option case.¹⁴ Agnes knows that ten miners work together at the bottom of one of two mineshafts but she doesn't know which shaft they're working in now. It starts to rain. If she does nothing, both shafts will flood and all the miners will be killed. She can block the first shaft entirely, ensuring that there is enough air in the first shaft for all ten miners and that the second shaft will flood killing anyone working in it. She can block the second shaft entirely, ensuring that there is enough air in the second shaft for all ten miners and that the first shaft will flood killing anyone working in it. Her remaining option is to partially block both shafts. If she does this, a little water will rush into both shafts. In each of the shafts there will be enough air for nine miners, but one miner will be certain to drown in whichever shaft they're currently in. Many people say that Agnes ought to partially block both shafts. If this is right, this means that the fact that, say, they were in the first shaft doesn't constitute a decisive reason for Agnes to block that shaft completely. (It would, however, if it were known to Agnes and thus part of her evidence.) This seems to suggest that what Agnes ought to do is not determined by the totality of objective fact and the intuition seems to be vindicated by the reasons perspectivist suggestion that the reasons that determine what Agnes ought to do are facts that are part of her evidence.

A third rationale appeals to some (putative) platitudes about normative reasons and guidance.¹⁵ Some would say that one essential job of normative reasons is to provide an agent with guidance. Suppose, however, that normative reasons were not subject to any sort of epistemic constraint so that the reasons that determine what the agent ought to do, as in Broome's example, are facts that the agent couldn't have any inkling of until long after the time of action. To respond to such facts, the agent would have to take wild stabs in the dark, to make bizarre guesses, to get themselves to act against their best judgment, etc. They could not be guided by the facts that should provide them with guidance if they are reasons. But this, the reasons perspectivists say, is deeply troubling. Lord (2018) thus proposes that the reasons that determine whether A ought to ϕ can only make it the case that A ought to ϕ in situations where A can follow their guidance in ϕ -ing by ϕ -ing for these very reasons. But an agent cannot ϕ for the reason that p , say, when they are in no position to know whether p , when they rationally believe $\sim p$, or when they lack the concepts necessary to even entertain the thought that p . That is to say, normative reasons are potent only when they have the potential to guide the agent by being the agent's reasons:

¹⁴ Parfit (2011).

¹⁵ See Kiesewetter (2016, 2017) and Lord (2018).

Normative Motivation: The reasons that determine whether A ought to ϕ are only facts that could be A's reason for ϕ -ing given A's epistemic position.¹⁶

Lord proposes that the Normative Motivation Thesis explains the Evidence-Relativity Thesis. Readers might disagree about the right explanatory order, but they might agree that there is something right about the underlying idea here about guidance and why that might motivate the Evidence-Relativity Thesis.

We will have more to say about the use of three-options cases in §3.2 and the Normative Motivation Thesis in §5. What we shall now do is explain why we think we should reject the proposal that evidence-relative reasons are normative:

Normativity of Evidence-Relative Reasons: A ought to ϕ iff there is a fact (or some facts) in A's evidence that constitute(s) a decisive reason for A to ϕ .

If the reasons perspectivists are right that, say, three-options cases show that facts external to our evidence are not normative, we should be able to conclude that normative reasons in general are not normative.

3. 1000 Children

We think that a simple example shows that the requirements of rationality differ from the demands generated by evidence-relative reasons. If so, we cannot appeal to the Normativity of Evidence-Relative Reasons to vindicate the Unification Thesis or defend the idea that rationality is normative. In §3.1, we shall present the example and our argument against the normativity of evidence-relative reasons. In §3.2, we show that we can modify our example so that it's a three-option case that threatens the normativity of evidence-relative reasons.

3.1 Punish the Children

Step 1: The base case.

Agnes has 1000 children. She knows each of their names (i.e., 1, 2, 3, ... Inge, ..., and 1000). When one of her children is reasonably suspected of having done something wrong, she presides over a trial and makes a decision about whether to punish or not. (Punished children are sent to their rooms, assigned extra chores, etc.) She uses the right evidence in such trials. (Readers should decide what that this condition comes to.) During one particularly trying weekend, each of her children

¹⁶ See Lord (2018). Kiesewetter (2016) defends a similar proposal. For a helpful survey of recent discussion of normative reasons and guidance, see Way and Whiting (2016).

seems to have broken some serious rule. She holds a trial for each child. After carefully reflecting on the evidence in each case, she concludes in each trial that the relevant child is guilty. The evidence is good enough for her to know in each case that the child is guilty. She concludes in each that the child ought to be punished. In reaching this verdict, she relies on this naïve reasoning: *this child is guilty. If so, I ought to punish. So, I ought to punish.* She acts accordingly.

Step 2: The duplicate.

Agnes* is Agnes' counterfactually situated internal mental duplicate (i.e., they are in precisely the same non-factive mental states, the mental events they undergo are subjectively indistinguishable, etc.). Because of this, Agnes* believes each of the things that Agnes knows. In particular, she has 1000 beliefs about the guilt of the children (i.e., that 1 is guilty, that 2 is guilty, ..., that Inge is guilty, ..., and that 1000 is guilty) and 1000 normative beliefs (e.g., I ought to punish 1, I ought to punish 2, ..., I ought to punish Inge, ..., and I ought to punish 1000). Since Agnes' 1000 beliefs about the guilt of the children constituted knowledge, Agnes*' 1000 beliefs about the guilt of the children are each rationally held. (This follows from our description of Step 1 and mentalism.) Since Agnes' 1000 beliefs about what she ought to do at the end of each trial are rationally held, the same is true of Agnes*' 1000 normative beliefs. Agnes and Agnes* punish each child in keeping with their normative beliefs.

There is one important difference between the two scenarios. Inge is innocent in the second scenario, but not the first. There is precisely one innocent child in the second scenario. Because of this, Agnes* has precisely one mistaken belief about the guilt of the children and Agnes has no mistaken beliefs about the guilt of any particular child.

Step 3: The informant.

A reliable testimonial informant tells Agnes and Agnes* that precisely one child was innocent. Agnes and Agnes* believe the informant. Alas, the informant disappears before the informant can say anything further. Agnes and Agnes* believe their informant, but they do not abandon their non-normative beliefs about the guilt of the children. Thus, they each have inconsistent beliefs about the guilt of the children (i.e., they still believe each child to be guilty but they also believe that not all the children are guilty). They face a practical question as to whether to continue the punishment or to free the children.

Step 4: The epistemic evaluation.

Agnes has 1000 beliefs concerning the guilt of individual children, 1000 normative beliefs about how to deal with them, and the belief that precisely one of these children is not guilty. Since she knew initially that each child was guilty, the belief that precisely one child is not guilty is mistaken. How many of her initial beliefs can retain their status as knowledge when she is told and comes to believe that precisely one child is innocent?

We think, optimistically, that it would be possible for each of these 1000 beliefs to retain their status as knowledge even after she believes her informant and forms the (mistaken) belief that one child is innocent. Not one of these beliefs is directly threatened by the informant's claim. We do not think it plausible that an arbitrary set of them must lose its status as knowledge nor that all must. (We'll discuss other options in due course.) So, we think that Agnes can still know of each child that the relevant child is guilty.¹⁷ And if Agnes knows each child to be guilty, she is rational in each of her beliefs about the guilt of the particular children and rational to believe in each case that she ought to punish the child. We see no reason to think that the naïve reasoning that she relies on would fail to make her relevant normative beliefs rational. Indeed, we think that she could know in each case that she ought to punish the relevant child.

What about Agnes*? In her case, we may assume, the informant knew that exactly one child was innocent. As before, we don't think this testimony has to destroy any of the knowledge that she had initially (e.g., that 1 was guilty, that 2 was guilty, etc.). What about the testimony itself? Could she come to know on the basis of the informant's testimony that one child was innocent? We think so. Just like authors often come to know that their latest book contains an error because a very reliable fact checker correctly testifies that this is so, we think that Agnes* can come to know that precisely one child was innocent. If this is right, Agnes* would know in all but one case that the relevant child was guilty, know that precisely one child is innocent, and rationally believe each child to be guilty. (This last claim follows from mentalism and our description of Agnes' epistemic situation. Whatever Agnes knows, Agnes* knows and/or rationally believes.) Agnes and Agnes* have the same 1001 non-normative beliefs, 1000 pieces of non-normative knowledge, but slightly different beliefs that constitute knowledge. As for Agnes* normative beliefs, we think that

¹⁷ The case is a variant of Makinson's (1965) preface case. For arguments that a rational agent can hold an inconsistent set of beliefs in such cases, see Christensen (2004), Easwaran and Fitelson (2015), and Worsnip (2016). For a dissenting view, see Ryan (1991, 1996) and Pollock (1986). The only discussion we've seen of knowledge in these settings is Moeller (2015) and we think that he's rightly sceptical of the idea that the testimony must destroy some of Agnes' knowledge. His discussion is primarily concerned with closure and the KK-principle, not the problems that rational inconsistency poses for reasons-first approaches to ought and rationality. For what it is worth, we do not think that Agnes can come to know that the informant is mistaken and we do not think that Agnes* can come to know that Inge is innocent. Because of this, we have to deny multi-premise closure.

in each case she would know and/or rationally believe that she ought to punish the relevant child. We can support this by appeal to mentalism and our description of Agnes' epistemic situation.

If this reasoning is sound, we have shown that the reasons perspectivists ought to accept the following:

Believe Each to be Guilty: Agnes and Agnes* ought to believe each child to be guilty.

Believe Each Ought to be Punished: Agnes and Agnes* ought to believe in each case that they ought to punish the relevant child.

Step 5: The practical evaluation.

Let's now think about Agnes' and Agnes*' practical situation. Remember that they accept the informant's testimony. We think that it is intuitively plausible that Agnes and Agnes* ought to punish most or all of the children. We think that this is independently plausible and in keeping with the reasons perspectivists' view:

Punish Each: Agnes and Agnes* ought to punish each of the children.

Consider the rational status of Agnes' and Agnes*' beliefs. Agnes considers the first child. Since she knows that this child is guilty, so she rationally believes it. She also rationally believes that she ought to punish this child if guilty. She rationally concludes that she ought to punish this child. Since it is irrational for her to violate the Enkratic Requirement, she cannot both rationally believe that she ought to punish and do other than punish, so she's rationally required to punish the first child. What holds in this case holds in the other cases. If, as reasons perspectivists say, rationality is normative, it follows from the fact that she is rationally required to punish in each case that she ought to punish in each case. If rationality is normative and mentalism is true, Agnes* ought to punish each because Agnes ought to punish each.

There is a second rationale for Punish Each, one that doesn't appeal to the Enkratic Requirement. In each case, Agnes knows that the particular child is guilty. Thus, in each case, she has an evidence-relative reason to punish, one that constitutes a decisive reason to punish. Thus, in each case she ought to punish.

What about Agnes*? She is Agnes' internal duplicate. The rational status of Agnes*' actions and attitudes cannot differ from the rational status of Agnes' actions and attitudes. If rationality is normative and the Unification Thesis is correct, Agnes* ought to punish when Agnes is rationally required to and Agnes* ought to believe when Agnes is rationally required to. So, Agnes* ought to believe each child to be guilty, ought to believe that each child ought to be punished, and ought

to punish each child. We don't need the Enkratic Requirement to justify Punish Each. We can use mentalism and the normativity of rationality to generalise from Agnes' case to Agnes*' case.

We claim that Agnes and Agnes* ought to believe most or all of the children are guilty, that they ought to believe that most or all of the children ought to be punished, and that they ought to punish most or all of the children. If these verdicts agree with the verdicts of reasons perspectivism, what could our objection be? The problem is this. Although reasons perspectivism gives us a rationale for Punish Each, it *also* gives us a rationale for rejecting Punish Each.

Focus on Agnes*. She knows that one child is innocent. She doesn't know which child, but because she knows that one child is innocent she knows that there is one case in which she is not aware of a fact or some facts that would constitute a sufficient reason for punishing a child. We think that if someone lacks sufficient reason for punishing a child, they ought not punish a child. So, we think that reasons perspectivists have to concede that Agnes* can know that she would lack a sufficient reason to punish if reasons are, as they claim, facts that constitute evidence-relative reasons. And if Agnes* can know that there is some child that she ought not punish, there is some child that she ought not punish. Reasons perspectivism are committed to Punish Each, but they are also committed to the denial of Punish Each. On their view, Agnes* ought to punish each child but Agnes* ought not punish all the children.

This kind of practical inconsistency is a problem for reasons perspectivism.¹⁹ We think that *if* Agnes* ought not punish all the children, it cannot be true that she ought to punish each of them. If we say that she ought not punish each of them, we either have to deny that rationality is normative or deny that evidence-relative reasons are normative. Since the reasons perspectivist accepts both and tries to ground the requirements of rationality in the demands of evidence-relative reasons, we think we have found a case that shows that this cannot be done. We think that the right thing to say here might just be that Agnes and Agnes* ought to punish each of the children because it's rational in each case to punish the children. If so, an agent like Agnes* ought to do things that she knows she couldn't have a single reason to do.²⁰

Among other things, we think that this case should give us a reason to worry about Normative Motivation. In each case, it will seem to Agnes* that she is aware of some fact that constitutes a decisive reason to punish (i.e., the fact that the relevant child is guilty). However, she

¹⁹ Recall that one of the main motivations for reasons perspectivism is the idea that normative reasons have to be able to give an agent guidance. If the guidance they give is that they tell an agent that they ought to do each thing on a list and ought not do everything on the list, we don't think that they're doing well as guides.

²⁰ Assuming, of course, that reasons are things like facts that constitute pros or cons, entries in a ledger, etc.

knows that *if* she punishes each child there will be one case in which she will carry out this punishment without there being an evidence-relative reason that constitutes a good reason to punish. She will know that in each case it will be as if such a reason exists, but she knows that this appearance will be illusory in one case. If we were to say that in the one relevant case (i.e., the case involving Inge) that she ought not punish, we get a counterexample to the Evidence-Relative Reasons Thesis and Normative Motivation. We can say that there *is* a potent and decisive reason for Agnes* not to punish (i.e., the fact that Inge is innocent) but this would be one that wasn't part of Agnes* evidence and thus couldn't have been her reason for not punishing. If, however, we were to say that this is a case where she ought to punish, the right explanation as to why this is would *not* appeal to the facts that Agnes* seemed to know that she would identify as her reason for punishing because such things are not facts. If we go this route, we would have to deny the normativity of normative reasons. This would be the dreaded case where Agnes ought to do something when there is not a single fact that constitutes a normative reason for her to do the thing in question.

It might be worth noting the following. It is consistent with everything that we have said that facts about what Agnes* ought to do and what Agnes* ought to believe might supervene upon facts about what Agnes*' evidence is. That could be so regardless of whether Agnes* ought to punish Inge or ought not punish Inge. The problem is that Agnes* knowledge does not include knowledge of what is and what is not included in her evidence. Negative and positive introspection fail for knowledge and for being in a position to know. Because of this, it shouldn't be surprising that there could be pairs of agents like Agnes and Agnes* who have very similar internal mental lives and very similar sets of reasons who have sets of reasons that support very different things. At least, they would support very different things if the *elements* of the sets of reasons were the things that determined what Agnes and Agnes* ought to do or ought to believe. Readers might then wonder whether it might be worth exploring the idea that the body of evidence taken collectively might be the thing that reasons perspectivists should focus on in deciding what Agnes and Agnes* ought to do. Perhaps they should. But they should bear in mind the problems about negative and positive introspection. If Agnes*' evidence includes the fact that some large set of children are guilty and lacks the fact that Inge is guilty, this is *not* something that she is in a position to know and thus it does not seem that this conjunction of all and only what Agnes* knows could possibly meet the access requirements imposed by the Normative Motivation Thesis.

3.1 Punish, Free, or Pay?

Agnes and Agnes* were given two options. When they heard from their informant and then thought about what to do with some child, they knew that they could continue the punishment or free the child. What happens if our informant reappears? What if the informant promises to tell Agnes* which child is innocent but will only give her this information for a price?

Here is a natural take on the example. When Agnes* is told that one of the children is innocent, she is likely to be uncertain about whether this is so. She believes this, but she also believes that each of the children is guilty. She might reflect on her situation and think to herself that two things are possible: (a) she knows each of the children to be guilty or (b) she knows 999 of the children to be guilty and falsely believes one to be guilty. So, she would know that if (a) is true, she knows a fact that constitutes a decisive evidence-relative reason to punish in each case and that if (b) is true that she knows a fact that constitutes a decisive evidence-relative reason to punish in each case but one where she would be aware of no such reason.

Agnes* thinks about the guilt of 666. She realises that she has three options. She can decide to free 666 or continue the punishment of 666 without paying or she can pay the informant and then decide whether to free or punish. She knows that if (a) holds it is better to punish without paying. She knows that if (b) holds it is better to punish without paying if indeed she knows 666 is guilty and better to free 666 without paying if indeed she is not aware of 666's guilt. Thus, it seems that she can know that it would be better to decide to punish 666 without paying or free 666 without paying. She knows full well that the evidence-relative reasons that she possesses will make a decisive case for punishing without paying or a decisive case for freeing without her having to pay the informant. So, it seems that the reasons perspectivist view doesn't predict that she ought to pay the informant before deciding. But we think that it's clear in this case that she ought to pay her informant. We don't see any interesting difference between this case and Parfit's (2011) mineshaft case. In Parfit's case, the agent knows that they ought to do what they know couldn't be objectively best. In our case, our agent knows that they ought to pay the informant even if this ensures that she doesn't do what's most favoured by her evidence-relative reasons.²²

4. Responses

²² Suppose that Agnes* has non-zero credence in (a) and (b). Accordingly, she doesn't assign extremal credence to 666's being guilty. We can set a price such that, in the light of Agnes*'s credences, paying has higher expected value than punishing 666 without paying and higher expected value than freeing 666 without paying. In the limit, we can make the information free. If it is free, reasons perspectivists can say that the agent *may* take the information first. But they cannot say that they *ought* to gather the information, which we take to be the right verdict.

As we see it, there are two serious objections to the reasons perspectivists' proposal. The first is that the view seems to deliver inconsistent verdicts about the case. While the reasons perspectivist view seems to support the idea that Agnes and Agnes* ought to punish each of the children and believe each child to be guilty, it also seems to support the idea that Agnes* could know that she ought not punish all the children since she would know that there was a case in which she would not be aware of any fact or facts that constituted a sufficient reason to punish. We want a view that delivers the right verdicts and we think that that requires that it meets a practical consistency requirement. We also think that the informant should be paid and that when we bring the informant back into the story we have a three-option case that shows that we ought to do things that we know couldn't be best given our evidence-relative reasons.

We shall consider some responses to our example. The first three responses might seem to give the reasons perspectivist a way to deal with our example. We think that they will not. They either don't deliver the right verdicts (i.e., they conflict with Punish Each or Believe Each to be Guilty, they miss the point of the example, or they introduce some proposals about reasons that we ought to reject). The next two responses require us to modify or abandon the reasons perspectivist view. These proposals suffer from a number of problems, but we think that they point us in the right direction.

4.1 Defeat

How might the reasons perspectivist respond? To generate the practical inconsistency, we argued that reasons perspectivists have to say that (a) Agnes* ought to punish each of the children since Agnes ought to and (b) that Agnes* ought not punish all the children because she knows that she is not aware of a sufficient evidence-relative reason to punish in one case. In arguing for (a), we had to assume that Agnes could know and rationally believe that each of the children were guilty even after the informant tells her that one child is innocent. The *defeatist* denies this. The defeatist doesn't think that Agnes' knowledge and rational belief could persist because they think that the informant's testimony defeats knowledge and rationality:

Defeat: It would be impossible for someone in Agnes' situation to know and to rationally believe each of the children to be guilty once the informant tells her that one child is innocent.²⁴

²⁴ The defeatist label might be misleading in this context. It was inspired by an idea that we find in Pollock's (1986) discussion of the lottery and preface paradoxes. He proposes that a kind of collective defeat ensures that our rational beliefs are consistent. Ryan (1991, 1996) defends a similar idea, but doesn't use the language of 'defeat'. Smith (2016, forthcoming) and Leitgeb (2017) defend views of evidential support that support the idea that rational and/or justified beliefs have

Since the continued possession of evidence-relative reasons requires the persistence of knowledge and of rationality, the defeatist thinks that they can undercut our objection reasons perspectivism by showing that neither Agnes nor Agnes* could possess reasons that required them to punish each of the children once one of them acquires the strong evidence of the false belief.

We think that we should distinguish between two defeatist views. The *strong* defeatist thinks that neither Agnes nor Agnes* could continue to rationally believe or know thanks to the informant's testimony. The *moderate* defeatist thinks that rational beliefs have to be consistent, but they don't endorse the stronger claim that neither Agnes nor Agnes* could continue to rationally believe that some of the children are guilty.

One problem with the strong response is its strength. Some of us are convinced by Christensen's (2004) suggestion that many of us have inconsistent beliefs, beliefs that could not all be true that do not contain any mistaken belief that we could identify as being particularly suspect. We believe many things but know that at least one of these things we believe is not true. This knowledge doesn't seem to prevent us from knowing a vast number of things. And if we are right that knowledge can be sufficiently robust to withstand the defeating power of the knowledge that some belief we hold but cannot identify is mistaken, we see no good reason to think that things would differ for Agnes and Agnes*.

A further problem for this defeatist proposal is that it seems to generate implausible practical consequences. If the strong defeatist response is correct, Agnes cannot rationally believe or know, say, that the first child is guilty. If she cannot rationally believe the child to be guilty, can she punish the child? We think that it is only proper to punish when it is proper to blame. We agree with Adler (2002) and Buchak (2014) that it is not appropriate to blame an individual for doing something if it's not rational for you to believe that they did it. Blame requires belief and conviction, not some weaker state. Adler's observation about resentment seems to hold true for blame as well:

Mild resentment is never resentment caused by what one
judges to be a serious offense directed toward oneself

to be consistent but they don't appeal to defeaters to explain why consistency is rationally required. We do not know what Leitgeb thinks about using statistical evidence for punishment and/or blame, but Smith is opposed to it and his view provides a neat and tidy explanation as to why it shouldn't be used in criminal trials. One worry about his framework, though, is that it seems to us to deliver the wrong verdict in preface-type cases. If he thinks that Agnes and Agnes* can rightly punish, it seems that he has to deny that it's only right to punish when it's rational to believe that the person being punished has done something wrong. And if he drops that, it's less clear why it should matter in the case of statistical evidence why (allegedly) we cannot have a rational belief in the guilt of the defendant.

tempered by one's degree of uncertainty in that judgment
(2002: 217).²⁵

Just as a high degree of confidence that fails to amount to conviction is not enough for resentment, it is not enough for blame. The Adler-Buchak point is that it is only proper to blame and punish when it is proper to believe that the relevant agent is guilty. The strong defeatist faces the problem that their view seems to support the conclusion that Agnes and Agnes* should not punish any of the children. Since neither Agnes nor Agnes* could rationally believe any particular child to be guilty, they could not reasonably blame any child or punish them for something that they cannot be blamed for doing.

We think that this is an overreaction. It restores practical consistency, but at too high a cost. We don't think that the discovery that a mistake was made should require us to open all the doors to the cells and empty our prisons. Some prisoners must be punished. And what goes for prisoners goes for children. Since some children must be punished. Punishment requires belief. Rational belief must be more resilient than the strong defeatist believes.

Of course, the strong defeatist can say that Agnes and Agnes* still know that most of the children were guilty. Unfortunately, this cannot give Agnes and Agnes* the right to punish any child in particular. Imagine Agnes' friend discovers that she punished her children in the knowledge that she didn't know that the child was guilty. If she said, 'Well, I had to punish Inge because most of the children did something wrong', we think that Agnes would show herself to be an unreasonable person. But we wanted to stipulate that Agnes was a reasonable person.

The moderate defeatist defends a more moderate proposal, which is that some of Agnes' and Agnes*' knowledge and rational belief succumbs to defeat. Does this view fare any better? We see two ways that the moderate defeatist view might be developed. The moderate defeatist might propose that in each case some *specific* set of beliefs lose their positive epistemic standings (i.e., some specific beliefs are longer rational to hold and no longer constitute knowledge). She might instead say that some *unspecific* set of beliefs lose their positive epistemic standings (i.e., while no rational inconsistency is possible, the defeat requires the agent to drop some of her beliefs without requiring her to drop any particular beliefs).

We don't think that there is any good reason to think that some specific belief or beliefs would have to lose their rational standing or lose their status as knowledge because the defeat doesn't target any beliefs in particular. There might be no discernible difference in, say, the quality of the grounds that support believing 666 is guilty from those that support believing 665 is guilty.

²⁵ For further discussion of the connections between emotion, reactive attitudes, and knowledge, see Dietz (forthcoming) and Unger (1975).

If the moderate defeatist believes that some specific belief is more liable to losing its rational standing than others when there is no discernible difference in the grounds that support them, their view of rational belief seems objectionably externalist.

There is a further problem with the moderate defeatist view if it rests on the idea that some specific belief must lose its rational status. We see no good reason to believe that the beliefs of internal duplicates will be equally likely to lose their status as knowledge. If, say, Agnes and Agnes* lose knowledge and thereby lose some reasons (thus ensuring that practical consistency is regained) but they end up losing different knowledge and thereby possessing different sets of evidence-relative reasons, the reasons perspectivist would have to say that they ought to punish different children. But this is incompatible with the combination of mentalism about rationality and the normativity of rationality.

Perhaps the best option for the moderate defeatist is one that says that rational beliefs are always consistent even though no particular belief or beliefs will be more liable to defeat than others. A rational thinker would have to maintain consistency amongst her beliefs, but she might be free to ‘choose’ which beliefs to drop to ensure consistency when the grounds for each belief are not discernibly different. Agnes would be rationally required to suspend on one or more beliefs about the guilt of the children and the same would hold for Agnes*. If the idea is that the agent has ‘free choice’ between different consistent subsets of propositions, then this moderate defeatist proposal tells us that the agent might rightly punish some children. The problem, as we see it, is that it seems arbitrary at this point for the agent to believe that some children rather than others are guilty. But then it seems arbitrary that some children are punished when others are not. Agnes runs a reasonable household because she’s a reasonable person. She wouldn’t make these arbitrary decisions between cases.

4.2 Objective vs. Non-Objective ‘Ought’

We anticipate a kind of knee-jerk reaction to our objection to reasons perspectivism, which is to appeal to some distinction between an objective reading of ‘ought’ and some other reading of ‘ought’.²⁶ We claimed that the reasons perspectivist has to say that (a) Agnes and Agnes* ought to punish each of the children and (b) Agnes* ought not punish all of the children. Someone might try to defend reasons perspectivism by claiming that the ‘ought’ in (a) is not the objectivist ought whilst claiming that the ‘ought’ in (b) is.

²⁶ For helpful discussions of objective and subjective readings of ‘ought’, see Olsen (2017), Smith (2010), and Wedgwood (2016).

The response will not help. Remember that we used the reasons perspectivists' commitments to justify (a) and (b). If, as we've argued, the defeatist response isn't convincing, we can assume that Agnes and Agnes* have 1001 beliefs that concern the guilt of the children where they each know 1000 things about their guilt. Because Agnes and Agnes* are internally the same and Agnes has the knowledge she does, we've shown that the reasons perspectivist has to accept (a). Because Agnes* has the knowledge that she does, the reasons perspectivist has to accept (b). Because we've argued from *their* assumptions and some independently plausible assumptions about knowledge, we cannot say that different readings of 'ought' are at play in (a) or (b) unless the reasons perspectivists' proposals are concerned with different readings of 'ought'.²⁷²⁸

4.3 Backup Reasons

Reasons perspectivists will most likely agree that (a) Agnes and Agnes* ought to punish each of the children but they will deny (b) Agnes* ought not punish all of the children. In this way, they shall try to find a way to restore practical consistency whilst finding a way to justify punishing all the children.

Recall that our justification for (b) was something like this. The reasons perspectivist should say that if Agnes* is not aware of a single reason to punish a child then Agnes* ought not punish this child. (Punishing someone is something we should not do unless there is sufficient reason to do it.) We said that because Agnes* knew that one child was not guilty, Agnes* knew that there was one case in which she was not aware of a fact that constituted an evidence-relative reason to punish the relevant child. The reasons perspectivist will contest just this point. They'll agree that Agnes* knew that there was a case in which she couldn't have been aware of the fact that the relevant child was guilty, but they will insist that Agnes* was nevertheless aware of a fact

²⁷ Some readers might still think that this 'ought' is too objective or externalist. We won't argue that this worry is misplaced, but we should stress that if this is right, this is a problem for the reasons perspectivist. We used the claims that they appeal to in order to defend the Unification Thesis to justify (a) and (b). If these claims concern different readings of 'ought' or a reading that is too objective or externalist, their defence of the Unification Thesis fails.

²⁸ We said earlier that our cases cause trouble for subjective reasons approaches to rationality. Our case suggests that Schroeder's (2008) proposal would need refinement since he identifies subjective reasons with reasons we would have if our relevant beliefs were correct. There is no possible world in which Agnes' and Agnes*' relevant beliefs would be correct because they are inconsistent. Thus, we don't see how we can straightforwardly apply this proposal to account for the intuition that it's rational for Agnes and for Agnes* to believe that Inge is guilty or to punish her. We think the case causes trouble for Whiting's (2014) proposal, too. On his proposal, Agnes has a subjective reason to ϕ only if it is apriori that if the facts are as they appear to her to be, those facts give her an objective reason to ϕ . Since 'the facts' cannot possibly be as they appear to her to be, we do not see how to apply this account to reach the right verdict.

that constituted a sufficient reason for Agnes* to punish Inge (i.e., the one innocent child). Our objection fails, they'll say, because we assumed that Agnes* would have to be aware of the fact that Inge was guilty. They'll say that Agnes*' awareness of some other fact could suffice.

They'll say that we've overlooked the possibility of backup reasons:

Backup Reasons: If A and A* are internal duplicates, A has decisive/sufficient evidence-relative reasons to ϕ iff A* has decisive/sufficient evidence-relative reasons to ϕ .

Agnes and Agnes* are internal duplicates but they know different things so they have different evidence-relative reasons. Still, the reasons perspectivists will say that their different reasons will support, require, favour, etc. the same things. If Agnes knows something that constitutes a decisive reason for her to punish Inge (when she is guilty), Agnes* will know something else that constitutes a decisive reason for her to punish Inge (even when she is innocent).

An example might be helpful. Gertrude orders a gin and tonic, is served a gin and tonic, and takes the first sip of her drink. She did so knowing that her drink was a gin and tonic. Because she's in the good case and we're not worrying about whether she's had too much to drink, we might say that she ought to drink, it's rational for her to drink, and that she has a sufficient evidence-relative reason to drink—the fact that the drink is a gin and tonic. Bernard orders a gin and tonic, but is served a Bernie (i.e., a petrol and tonic). Because Bernard is Gertrude's internal duplicate, we might think that it's just as rational for him to drink as it was for Gertrude to drink. In the reasons perspectivist framework, this could only be true if Bernard is also aware of some facts that make a sufficient case for drinking. What would Bernard's evidence-relative reasons be? They would be things he knew, such as the fact that it seemed that it was a gin and tonic, the fact that the bar was reliably run, the fact that the glass contained ice, tonic, and limes, etc.²⁹ The Backup Reasons proposal requires us to find some fact or facts that could in the bad case make a sufficient or decisive case for doing the things that there is sufficient or decisive reason to do in the good.

Applied to our case, we would have to say that in each case, Agnes* is aware of some fact or facts that make a decisive or sufficient case for punishing. In the case where guilt is known, the fact is the guilt of the child. In the case where the child is innocent, the fact is something else (e.g., that the child appears guilty, that there were certain clues that indicated guilt, etc.).

²⁹ Kiesewetter (2017) and Lord (2018) defend this sort of response to error cases. See also Ichikawa (2018).

We do not think that the Backup Reasons proposal works for things like blame, resentment, or anger. Consider Alvarez's example.³¹ 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 good reason for Othello to be upset, the fact that it appeared to Othello that Desdemona was unfaithful was an *equally good* reason for Othello to be upset with her.

This cannot be right. It is no part of Shakespeare's play that Othello would have thought that a good reason for being upset could have been that it appeared to him that Desdemona was unfaithful. 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). I 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.³² We should not generally assume that if the fact that p constitutes a decisive reason for an agent who knows p to ϕ , 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 ϕ . 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.

³¹ Alvarez (forthcoming).

³² 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.

4.4 Alternative Facts

We have run out of options for reasons perspectivists. We haven't found a way to reconcile our intuitions about our case with mentalism about rationality and the idea that evidence-relative reasons are normative. We would like, however, to have a coherent treatment of our case. Let's look for an alternative to reasons perspectivism.

There is a new view that shares the reasons perspectivists' aim of trying to account for ought and rationality in terms of evidence-relative reasons. The main difference between this view and reasons perspectivism is that this new view denies factualism:

Alternative Facts: The reasons that justify an agent's actions and attitudes are the propositions that it is rational for the agent to believe.³³

On this view, internal duplicates are rational to respond in the same ways and ought to do the same things because the relevant normative reasons are just the propositions (true or false) that capture how things seem or appear to the agent and/or are justifiably/rationally believed by the agent. On this view, Agnes and Agnes* can be aware of the same reasons. In each case, the agent is aware of a decisive reason to punish (i.e., the proposition that the relevant child is guilty). On this view, we have no reason to think that Agnes* can come to know (by drawing an inference from the informant's testimony) that she ought not punish one of the children because the information that the informant gives her would not reveal that there was a case where she was not aware of a decisive reason to punish. When Agnes punishes Inge, she does so in the knowledge that Inge was guilty. When Agnes* punishes Inge, she does so while rationally believing that Inge is guilty. The knowledge and the rational belief is sufficient for awareness of the reason that constitutes a decisive reason to punish (i.e., the proposition that Inge is guilty). This view ensures practical consistency because it undercuts our argument for thinking that Agnes* can come to know that she lacks sufficient reason to punish one of the children when she learns that one child is innocent. If she knows that in each case she has a rational belief in guilt, she would know that in each case she would have sufficient reason to punish the child even if the child happened to be guilty.

³³ See Comesaña (forthcoming), Comesaña and McGrath (2014), Drake (2018), and Fantl and McGrath (2009). We think that the linguistic arguments in Unger (1975) and others support the idea that propositionally specified ascriptions of reasons are factive, but this might be taken to be a point about language and there are perfectly good ways of stating the view in question without using our ordinary ways of speaking.

The backup reasons view says that Agnes* had sufficient reason to punish Inge (i.e., the fact that Inge appeared guilty, say) and the alternative facts proposal agrees that Agnes* had sufficient reason but offers a different account of what that reason is. Notice that the view will differ from the defeatist view only if it allows that Agnes and Agnes* can have rational but inconsistent beliefs about the guilt of the children (i.e., that 1 is guilty, that 2 is guilty, ..., that 1000 is guilty, and that one child is not guilty). Since we're looking for alternatives, let's suppose that the proponents of alternative facts want to say that Agnes and Agnes* to have rational but inconsistent beliefs.

The main problem with this approach is that it seems to only allow for rational inconsistency among the agent's beliefs by recognizing the possibility of inconsistent normative reasons. We think that normative reasons should be consistent. What's wrong with inconsistent normative reasons? A number of authors, including those who defend the alternative facts proposal, defend the view that the agent's reasons are things that the agent can rationally presuppose in shaping the decision matrices that the agent ought to follow (Comesaña forthcoming).³⁴ In our example, if Agnes and Agnes* justifiably believe, say, that Inge is guilty, they can presuppose her guilt in formulating the decision matrix that they ought to follow and think about whether, in light of this guilt, it would be best to punish. They can also presuppose that 1 is guilty, that 2 is guilty, that 3 is guilty, and so on. They can also presuppose that one is not guilty. We see no coherent way to construct a decision matrix that, say, presupposes the guilt of each whilst presupposing the innocence of one. The matrix is supposed to represent states, options, and the value of the outcomes (i.e., state-option pairs) and the states that we're supposed to represent are the ones compatible with the Agnes' and Agnes*' justified or rational beliefs. But there are no states compatible with these beliefs. We don't know what it would mean to say that Agnes and Agnes* ought to follow the recommendations of the decision matrix if we identify their evidence-relative reasons/normative reasons with the inconsistent propositions that they rationally or justifiably believe.³⁵ Until it is developed further and represents a genuine alternative proposal

³⁴ This is also the view of Littlejohn (2012), but he thinks that justified beliefs have to be true. Dutant (forthcoming) and Weatherson (2012) have proposed that the condition under which we might presuppose belief is when the belief constitutes knowledge.

³⁵ Comesaña (forthcoming) acknowledges that in such cases, Bayesian decision theory can no longer be applied to tell us what Agnes and Agnes* ought to do or what would be rational for them to do, but he offers no alternative. Thus, we fear that while the introduction of alternative facts might have some virtues (e.g., it helps us vindicate mentalist intuitions without the need to introduce backup reasons), it undermines the best stories we have about why Agnes and Agnes* ought to punish the children they know to be guilty. Is this fatal? We do not think so. Comesaña could say that the things we should presuppose in formulating a decision matrix are *certainties*. This might come with sceptical costs since it forces him to say that we're only rational or justified in

to standard decision-theory, we do not see the appeal to alternative facts as a useful move for explaining what Agnes and Agnes* ought to believe, feel, and do.³⁶

4.5 Probable Guilt and Practical Adequacy

We have looked but have not found a normatively adequate treatment of our case. Each view had this in common. Each view has tried identify some fact or proposition that could constitute one of Agnes and Agnes*' evidence-relative reasons and be a decisive reason for Agnes and Agnes* to punish each of the children. One proposal found plausible candidates for playing this role but it told Agnes* that she ought to punish each of the children and ought not punish all of them. The other proposals did better on the practical consistency score but either told her not to punish the children or offered some problematic account of what constituted a decisive reason for Agnes* to punish innocent Inge. Perhaps we should turn our attention towards a view that isn't reason-centric in the way that these views have been. Not everyone seems to agree that normative reasons are normative.³⁷

Here is a new proposal. Agnes and Agnes* ought to punish each of the children *not* because they are aware of a fact or a false proposition that constitutes a decisive reason to punish innocent

believing when we're rational or justified in being completely certain, but it might help him avoid having to say that there can be inconsistent sets of normative reasons. He might then appeal to one of the proposals below to address the problematic cases involving rational but inconsistent beliefs.

³⁶ We might add one further worry. If recognizing the possibility of rational but inconsistent beliefs turns on whether we can recognize the possibility of inconsistent evidence, we run into an epistemological problem. To handle our case in a way that is consistent with mentalism, we want a view on which Agnes and Agnes* can rationally believe each child to be guilty whilst believing that one is innocent. (Without this, it seems we have to say that Agnes and Agnes* ought to free each of the children, but this is an overreaction to the (apparent or real) discovery of one innocent child.) Suppose that Agnes' evidence includes each of the propositions that she rationally believes (i.e., that 1 is guilty, that 2 is guilty, ..., that 1000 is guilty, and that one is innocent). What is the evidential probability on this view that 1 is innocent? It would seem to be undefined. The evidential probability that 1 is innocent is the conditional probability of this proposition on Agnes' total evidence, but given the standard definition of conditional probability, this would be undefined when the probability of Agnes' total evidence is itself 0. This is a point that Williamson (2009) presses in response to Goldman's (2009) suggestion that the contents of a thinker's justified, false beliefs belong to that thinker's evidence. As with the problem of understanding how to formulate a decision matrix that presupposes an inconsistent set of propositions, we do not think that such problems are fatal. We think that they are challenges that a view must address in order to be considered as sufficiently developed alternatives to extant theoretical options that seem to suffer from none of these problems and seem perfectly well-equipped to handle the cases we throw at them.

³⁷ Readers should check Zimmerman's (2008) subject index to see if they can find an entry for 'reason'.

children but because Agnes and Agnes* need to punish each child in order to maximise expected value:

Probable Guilt: Agnes and Agnes* ought to punish a child iff
the probability of the child's guilt is sufficiently high.³⁸

On this approach, the presence or absence of evidence-relative reasons to punish is irrelevant to the determination of what Agnes and Agnes* ought to do. There is, if you like, a principle of rationality that enjoins them to act, one that tells them what they ought to do or what is rationally required of them given the values at play and their evidence.

Think about the values that we would plug into this decision-matrix:

	Guilty	Otherwise
Punish	a	b
Don't Punish	b	d

We assume that a is at least as good as d, and that d is better than b and c. Unless we think that Agnes and Agnes* ought to be pathologically averse to the possibility of punishing an innocent child, we should see that the expected value of punishing a child exceeds that of not punishing this child. So, Agnes and Agnes* ought to punish each child.³⁹ Since Agnes and Agnes* should know that it doesn't matter whether they are actually aware of the guilt of the child in any particular case, Agnes* cannot infer that she ought not punish each child when she learns that one child is innocent.

Here is a natural worry about this proposal. Although it rightly sanctions punishing loads of children, it also sanctions punishing children when we don't believe that the relevant children have done anything wrong. Here is a way to try to address this worry. Just as probability of guilt might help us decide whether Agnes and Agnes* ought to punish, it can help us decide whether

³⁸ See Lempert (1977) and Steele (MS) for a defence of this proposal in a legal setting. We shouldn't assume that they think Agnes should treat her children like criminals. Their view is attractive if we frame the decision-problem in the way that they have and think that it's rational to do that which maximises expected value. For reasons discussed below, we don't think that the problem should be framed in the way that they have.

³⁹ Instead of thinking of this in terms of 1000 individual choices between two options, we could think of this as one choice with 2^{1000} options (e.g., punish all, punish none, punish all but the first, punish all but the second, punish all but the first and second, etc.). The expected value of punish all (i.e., $1000 * (\text{Pr}(g) * u(a) + \text{Pr}(\sim g) * u(b))$) exceeds the expected value of punish none ($1000 * (\text{Pr}(g) * u(c) + \text{Pr}(\sim g) * u(d))$). The expected value of punishing some is greatest when all children are punished (i.e., the value of $(1000-X) * (\text{Pr}(g) * u(a) + \text{Pr}(\sim g) * u(b)) + (X) * \text{Pr}(g) * u(c) + \text{Pr}(\sim g) * u(d)$ is greatest when $X = 0$).

Agnes and Agnes* ought to believe that a child is guilty. If we adopt a Lockean view of rational belief and we assume that the probability of guilt is sufficiently high, it might be possible to defend a view on which the minimum probability of guilt required for permissible punishment is sufficiently high to make it rational for Agnes and Agnes* to believe in each of the relevant cases that the child is guilty.⁴⁰

One nice thing about this proposal, then, is that it seems that it can accommodate the Adler-Buchak point. Relatedly, it might seem that this proposal also vindicates the naïve reasoning that Agnes and Agnes* rely on. In each case where Probable Guilt says that they ought to punish, the Lockean might also say that they rationally believe that the child is guilty. And if they rationally believe the child is guilty, they might be in a position to rationally infer that the child ought to be punished.

Alas, we think that there is a problem with this proposal. In fact, we think there are two closely related problems with this proposal. First, we think that there is a kind of naïve reasoning that we rightly rely on when punishing (i.e., *This child is guilty. If so, I ought to punish. So, I ought to punish*) and we want an account that tells us that when it's not rational to rely on such reasoning, it's not right or rational to punish. (This is connected to the Adler-Buchak point.) There is a second kind of naïve reasoning that we rely on when deciding not to punish: *I don't know if this child is guilty. If so, I ought not punish. So, I ought not punish*. We also want a view on which it's not rational or right to punish when it's rational to rely on this second piece of naïve reasoning. We don't think that the Probable Guilt view will get this right. We think there are cases where an agent will be very certain that a child is guilty and even more certain that they cannot know if that child is guilty. If they were to rely on both pieces of naïve reasoning, they would conclude that they ought to punish and ought not punish. Indeed, given certain assumptions about rational belief, it might be possible for an agent to believe *both* that a child is guilty and that they do not know if this child is guilty. We think this is a bad thing. Second, we think that there is a kind of case that shows that Probable Guilt delivers the wrong verdict. We think that this case illustrates the problem that the second kind of naïve reasoning poses for Probable Guilt when it is combined with the Lockean view.

Here is the problematic case:

Prisoners: 100 prisoners are exercising in the prison yard.
Suddenly 99 of them attack the guard, putting into action a plan that the 100th prisoner knew nothing about. The 100th prisoner played no role in the assault and could have done

⁴⁰ For helpful defences of the Lockean view, see Dorst (forthcoming), Easwaran (2016), Foley (2009), and Sturgeon (2008).

nothing to stop it. There is no further information that we can use to settle the question of any particular prisoner's involvement (Redmayne (2008)).

We think this case illustrates a number of things.

First, we think that it would be wrong to punish a prisoner taken from the yard. We think that this is connected to the fact that we know that we couldn't know whether they were involved. This seems to be a counterexample to Probable Guilt. If Probable Guilt gets our case right, it gets it right for the wrong reason.

Second, it suggests that there is something wrong with the decision-matrix that we sketched above. Focus on the outcome associated with the top-left cell. When we described the matrix, we described this as a positive outcome. It might be in some cases (e.g., Agnes punishes 666). It might not be in other cases (e.g., someone punishes the prisoner chosen from the yard). This suggests that we need a different matrix, one that recognises that value realised by punishing someone who is guilty might differ depending upon whether we know them to be guilty or know that we couldn't know.

Third, we think that it shows that naïve reasoning is a problem for a view that combines the Lockean view of rational belief with Probable Guilt. In Prisoners, the agent could, on the Lockean view, rationally believe the defendant to be guilty (owing to the high probability of guilt) *and* rationally believe that they don't know if the defendant is guilty (owing to the high probability that the guilt of the defendant is not known). So, while the view might vindicate the first kind of naïve reasoning, it doesn't vindicate the second.

Fourth, we think that the case might tell us something about rational belief. One nice thing about the Lockean view is that it allows for rational inconsistency. In turn, this allows us to say that Agnes and Agnes* can rationally believe each child to be guilty and thus never punish a child without fully believing the child to be guilty. Alas, we worry that just as probable guilt might not be sufficient for rightly and rationally punishing a child, probable truth might not be sufficient for rational belief. When the thinker knows that they cannot know p , the high probability of the truth of p isn't enough for rational belief. We shall exploit this fact in providing a better treatment of our case.⁴¹

⁴¹ We think that this case also poses a challenge to people who think that we can know lottery propositions (e.g., Sosa (2015)) or believe that reliability might be sufficient for justification (e.g., Goldman (1986)) since we think that these views, like the Lockean view, struggle to explain why we shouldn't blame or punish in cases like Prisoners. For reasons discussed by Harman (1968), we think that views that link justification and/or knowledge to reliability and/or probability struggle to get the lottery-type cases right. See Nelkin (2000) for further discussion.

These problems with Probable Guilt point to a problem with the pragmatic encroachment view. As it was developed by Fantl and McGrath (2009), the pragmatic encroachment view tied the justification of belief to the propriety of treating that belief's content as a reason for action. It seemed to provide a very nice framework for fallibilists (i.e., people who think that when we typically have justified beliefs about the external world, we rationally assign a probability less than 1 to their contents) who wanted a view on which justified beliefs are properly included in practical deliberation. On this view, a belief would be justified only if it would be proper to treat its content as a reason for action.

As we saw earlier, some link treating a belief as a reason for action as a matter to presupposing it for the purposes of practical reasoning. To determine whether this would be proper, the pragmatic encroachers have relied on a practical adequacy test:

Practical Adequacy: Agnes and Agnes* can properly treat the contents of their justified beliefs as reasons for action only if the action that would be recommended if they presupposed the truth of those beliefs is the same as the action that would be recommended given the probabilities that they rationally assign to these contents.⁴³

While this view seems to handle many cases well, we think that it does not handle the present case well at all. If we were to just focus on the probabilities that Agnes and Agnes* assign to the guilt and innocence of the children, we don't seem to have the right information to determine whether they ought to punish. We have rejected Probable Guilt, so we reject the idea that the probabilities of guilt and the values assigned to the decision matrix from above determine what the agent should do. Thus, it seems that the pragmatic encroachers either give us the wrong advice (i.e., to follow the guidance of the probable guilt view) or no advice (i.e., to do something else but without identifying any ground for preferring punishing to refraining from punishing).

We think this problem has been overlooked. It seems that to determine whether it would be right to punish, we first need to know whether it would be right to believe that a child is guilty and whether blame is appropriate. We need to use the propriety of belief and blame to determine the rightness of action. The pragmatic encroacher needs to use the rightness and/or rationality of the action to help determine whether belief and, by extension, blame is appropriate. Thus, we think that the view cannot give us any guidance in the case we have before us. It is one of the

⁴³ See Fantl and McGrath (2009) and Owens (2000) for defences of pragmatic encroachment views that seem to rely on Practical Adequacy. See Anderson and Hawthorne (forthcoming) for discussion of the connection between Practical Adequacy and pragmatic encroachment.

reasons that we find the present case interesting. We think it reveals a problem in the way that pragmatic encroachers approach their cases.

There is one further problem with an appeal to pragmatic encroachment in this setting. Recall the three-option variant of our example. An informant will, for a fee, identify any innocent child for Agnes and Agnes*. If the price is low enough, we saw that it seems rational to punish only after paying for the information. If this is right, the practical adequacy test tells us that neither Agnes nor Agnes* could rationally believe any of the children to be guilty until they pay the price. We think that this isn't a plausible result. As we've said, we think that Agnes and Agnes* can know and rationally believe that the children are guilty throughout the story. What's worse for Practical Adequacy is that it seems that Agnes and Agnes* could retain their rational beliefs and knowledge about the guilt of the children if the price were too high. We think that it's odd that it should be easier for Agnes and Agnes* to know if the informant were to ask too high a price for the information.

5. Knowledge and Probable Knowledge

We want a view that tells us that Agnes and Agnes* ought to punish most or all of the children, ought to only punish when it's rational to believe that a child is guilty, and that says that they ought never punish a child when they know that they don't know that the child is guilty. Here are two.

5.1 Known Guilt

This view does better on these points than the views considered thus far:

Known Guilt: Agnes and Agnes* ought to punish a child iff they know that some particular child is guilty and ought to believe that a child is guilty iff that belief constitutes knowledge.⁴⁴

If readers wanted a view that's compatible with the idea that normative reasons are normative, this view has that virtue. It tells us, for example, that Agnes* ought not punish Inge and ought not believe that Inge is guilty. It tells us that Agnes and Agnes ought to punish the rest while believing them to be guilty.

While this view is compatible with the idea that normative reasons are normative, it clearly clashes with mentalism or with the normativity of rationality. Moreover, it is incompatible with

⁴⁴ Blome-Tillmann (2017) defends a knowledge ideal for conviction in legal settings, though he derives a weaker norm from it. For defences of the idea that knowledge is necessary and sufficient for justified belief, see Sutton (2007).

the idea that evidence-relative normative reasons are normative. If normative reasons are normative and Agnes* ought not believe that Inge is guilty and ought not punish Inge, there must be something that constitutes the decisive reason for Agnes* not to punish Inge and something that constitutes the decisive reason for Agnes* not to believe that Inge is guilty. In neither case would the relevant reason be a fact that Agnes* knew or was in a position to know. The existence of such a reason supervenes upon Agnes* total evidence (on the assumption that Agnes*' total evidence is just her total knowledge), but that doesn't make the relevant fact or facts accessible to Agnes*. Remember that positive and negative introspection fail for knowledge and so would fail for evidence/reasons on the assumption that the facts that constitute a thinker's evidence or evidence-relative reasons are things known to the thinker. There will be things known that the thinker cannot know they know and things that the thinker fails to know that they cannot know they fail to know.

For these reasons, we suspect that most readers won't be thrilled with this view. It suffers from another potential defect. It isn't clear that what it says about the three-option case. Should Agnes and Agnes* pay the informant? If we assume that Agnes and Agnes*' evidence includes everything they know and we assume that they ought to proportion their beliefs to the evidence, we get the result that they shouldn't pay the informant. Agnes should be certain that each child is guilty and Agnes* should be certain that Inge is innocent. This probably doesn't sit well with most readers. If, on the other hand, we think that Agnes and Agnes* might be rationally uncertain of the guilt of a child even when they know that the child is guilty, we might get the result that Agnes and Agnes* should pay for the information, but this might not be enough to convince readers that this is the best possible view. We have saved that view for last.

5.2 Probable Knowledge

Let's consider one final view:

Probable Knowledge: Agnes and Agnes* ought to believe a child is guilty iff it is sufficiently probable that they know a child to be guilty and ought to punish a child iff it is sufficiently

probable that in so doing they punish a child they know to be guilty.⁴⁵⁴⁶

As with the previous view, this view delivers the right verdicts when it comes to belief, when it comes to punishment, when it comes to the links between belief and punishment, and when it comes to decisions about whether to pay informants.⁴⁷

This view is similar to Probable Guilt insofar as it says (roughly) that an agent ought to act and believe in such a way to maximise a kind of expected value, but it differs from that view in terms of what it takes the relevant values to be.⁴⁸ In the case of belief, it says that a rational believer's

⁴⁵ This view is similar to Blome-Tillman's (2017) proposal for legal settings and Dutant and Fitelson's (MS) proposal about rational belief. Blome-Tillman's norm is a narrow-scope norm: if it is probable that the court knows p , p should be admitted as evidence. Dutant and Fitelson (MS) propose a wide-scope rational requirement: one should believe p only when the thinker has high credence that she knows p . Both views differ from Moss' (2018) view, which links blame and punishment to various levels of probabilistic knowledge. We want to focus on punishing children instead of criminal law because we think the legal setting can introduce complications that distract from the issues that interest us here. One complication that we would like to avoid concerns the role that moral encroachment might play in explaining why we shouldn't punish using statistical evidence. Moral encroachment plays no role in our proposal, but seems to play an important role in Moss' proposal. For a critical discussion of moral encroachment, see Gardiner (2018). In terms of how it handles the difference between our preface-type case and the lottery-type case (i.e., Prisoners), it seems to have similar results as Neta's (2009) proposal about treating things as reasons for action.

⁴⁶ One interesting idea would be to combine this view and Known Guilt into a single view (e.g., by treating Known Guilt as a standard that determines objective obligation and by treating Probable Knowledge as a standard that determines rational ways of proceeding in light of uncertainty about what it takes to meet the objective standard). Investigating this further is beyond the scope of this discussion. It is ruled out by views that characterise the more subjective notion in terms of beliefs about the objective notion or in terms of counterpart relations between duplicates, but these might be non-optimal ways of understanding the relationship between the notions. For a helpful discussion of how to link the objective and subjective, see Jackson and Smith (2016).

⁴⁷ We do *not* think that the probability of knowing some proposition or propositions is essential to all rational choice. We think that our practices of holding people responsible, blaming them, and punishing them is one in which we place very different values on praising and blaming *because* someone has done something praiseworthy or blameworthy and doing such things when it happens to be the case that someone has done something praiseworthy and blameworthy. If this is right (and some further assumptions mentioned below hold), Probable Knowledge could be seen as following from the idea that we ought to act and believe in ways that maximise expected value. We think that the kind of act in question differs from, say, betting on football precisely in that the values realised by the outcomes in our cases are knowledge-sensitive in ways that bets on football matches, investments in the stock market, and plays at the roulette wheel are not.

⁴⁸ We do *not* think that the probability of knowing some proposition or propositions is essential to all rational choice. We think that our practices of holding people responsible, blaming them, and punishing them is one in which we place very different values on praising and blaming *because* someone has done something praiseworthy or blameworthy and doing such things when it happens to be the case that someone has done something praiseworthy and blameworthy. If this

beliefs should realise *gnostic* value. Whereas a *veritist* believes that the fundamental good is accuracy, the gnostic believes that it is knowledge. And while the defenders of Probable Guilt said that we ought to punish when the probability of guilt is sufficiently high, the defenders of Probable Knowledge stress the value of punishing those *known* to be guilty. Thus, like Probable Guilt, Probable Knowledge says that it doesn't matter whether Agnes or Agnes* is actually aware of the guilt of the child when they punish. It thus predicts that an agent like Agnes* might rightly punish a child such as Inge even when Agnes* is not aware of the guilt of the child. It thus denies that evidence-relative reasons are normative.

In explaining why Agnes ought to punish each child, the proponents of Probable Knowledge tell us that their decision-matrix has advantages over the veritist decision-matrix considered above.⁴⁹ Let's contrast the two matrices:

Veritist Decision-Matrix

	Guilty	Otherwise
Punish	+	-
Don't punish	-	0

is right (and some further assumptions mentioned below hold), Probable Knowledge could be seen as following from the idea that we ought to act and believe in ways that maximise expected value. We think that the kind of act in question differs from, say, betting on football precisely in that the values realised by the outcomes in our cases are knowledge-sensitive in ways that bets on football matches, investments in the stock market, and plays at the roulette wheel are not.

⁴⁹ For further discussion of these matrices and discussions of more complex options, see Dutant and Fitelson (MS) and Littlejohn (forthcoming b). Note that we treat punishing the guilty who are not known to be guilty (and believing the guilty to be guilty without knowing them to be guilty) in a column that realises a negative outcome. Is this plausible? We think so. In the case of full belief, we think that there are accurate or true beliefs that realise no epistemic good at all and indeed are bad states to be in (e.g., believing what one knows one couldn't know). In the case of punishment, we think that there are cases of punishing the guilty that bring about bad outcomes (e.g., punishing someone who happens to be guilty on the basis of reasons that one should have known would not put one in a position to know that this person is guilty). Part of our explanation as to why we think we should carve things up this way is that we think that the practice of punishing and blaming are regulated by certain norms or standards that we would not intentionally impose harms upon others unless that was done for certain reasons (e.g., facts that show that the person being harmed deserves to suffer the relevant harm). An act that fails to be based on such reasons, we say, involves a failure to properly use the power that comes with the authority to punish. So, we think of cases where someone is harmed as a consequence of someone's intention to punish where the harm is not suffered because of a wrong that merits this response is a bad outcome. And we think that an agent only harms someone because they have done wrong if the agent knows that the person being punished has done wrong. Similar points apply to praising and rewarding. Discussion of such cases would have made for a cheerier paper.

Gnostic Decision-Matrix

	Known to be Guilty	Otherwise
Punish	+	-
Don't punish	-	0

While the defenders of Probable Guilt and Probable Knowledge will agree that, say, Agnes ought to punish 666, they disagree about why this is. The defenders of Probable Guilt will say that it's because it maximises expected veritist value (i.e., it serves the values plugged into their matrix). The defenders of Probable Knowledge agree with the verdict but challenge the explanation. Probable Guilt tells us that an agent ought to punish if they're in Agnes' shoes and that they ought to do so for the very same reason that they ought to punish a prisoner in Prisoners. This is just what the defenders of Probable Knowledge object to. They do not see the outcome of Punish in Prisoners as positive. Indeed, they see it as a negative outcome—the relevant agent would punish someone while being certain (or highly confident) that they could not know that they were guilty.

Probable Knowledge has the virtue of vindicating both kinds of naïve reasoning. When it's rational for an agent to believe that a child is guilty, it's rational to punish. When it's rational to believe that the guilt of a child isn't known, it's not rational to punish. We thus vindicate the idea that Agnes rightly punishes only when she can rationally reason *as if* she is aware of a decisive reason to punish. We also vindicate the idea that Agnes rightly decides against punishing whenever she can rationally reason as if she's not aware of a sufficient reason to punish. Of course, *we* know that when she runs through this reasoning, the situations in which she rightly deploys this reasoning only imperfectly correlates with the cases where the premises of her reasoning are correct. Agnes* can rationally run through the naïve reasoning before punishing in each case, but the premises operative in her reasoning only correspond to the facts in 999 cases out of 1000. Thus, the view gives us practical consistency, a case for punishing loads of ill-behaved children, and it vindicates the Adler-Buchak point. After struggling to find a view that checks all these boxes, this is no mean feat.⁵⁰

⁵⁰ One interesting difference between Known Guilt and Probable Knowledge is that the former says that it would be right to punish in cases of improbable knowing and the latter says that it needn't be. (For discussion of improbable knowing, see Williamson (2017).) We think that this might be a virtue of the Probable Knowledge proposal. After all, we suspect that many readers would agree that if the probability that the guilt of the defendant is known is incredibly small it would be unreasonable to punish. On this view, it does mean that the mere fact that an agent knows that someone has done something wrong is not itself sufficient to ensure that they are warranted in blaming. In such a case, Probable Knowledge says that this individual wouldn't be rational in believing the first premises in the naïve reasoning for punishment or the naïve reasoning

Probable Knowledge is compatible with the idea that rationality is normative and mentalism about rationality. There is no case where the defenders of Probable Guilt have to say that Agnes ought to do something irrational or that she's rationally required to something that she ought not do. It allows us to say that Agnes and Agnes* ought to punish each of the children and ought to punish only those they can rationally believe to be guilty, provided that it's rational to believe what we probably know. The view ensures practical consistency just as Probable Guilt did. When Agnes* comes to know that one child is innocent, she does not come to know something that ensures that she ought not punish one of the children.

Like Probable Guilt, Probable Knowledge quite naturally explains why Agnes and Agnes* ought to pay an informant before punishing. When we expand the decision-matrix to include this additional option, the choice to punish only after paying maximises expected gnostic value. Since the introduction of this new option does not affect the probability that Agnes or Agnes* know, it doesn't generate the sceptical pressure associated with Practical Adequacy. The defenders of Probable Knowledge will say that there are situations where a rational thinker will believe someone to be guilty but seek further evidence before punishing, but we don't think that this is all that counterintuitive. It is much better for a view to say this than to say that it's right to punish those who we know we don't know to be guilty or those that we cannot rationally believe to have been guilty.

6. On the Normativity of Rationality and Reasons (Understood as Facts)

Let's suppose that Probable Knowledge and the gnostic value theory that underwrites it are correct. If so, it looks like much of the recent work on normative reasons and their connection to rationality and normativity turns out to be mistaken. Three-options cases seem to rule out any view that identifies normative reasons with facts that aren't part of an agent's evidence. Taking this much as read, the only remaining option would be a view that identifies normative reasons with something like facts that constitute evidence-relative reasons. But this view struggles with preface-type cases. In each case, Agnes* ought to punish. She ought to punish, in part, because she rationally believes each child she punishes is guilty. It would be a mistake to try to find in the contents of the beliefs that rationalise her decision to punish the fact or the facts that constitute the decisive reasons that make it right for her to punish. We might pair off the attitude and the fact when she punishes 1, punishes 2, ..., punishes 666, etc., but not when she punishes innocent

against punishment. In turn, this suggests that a kind of unreasonable knowledge is possible. See Lasonen-Aarnio (2010) for discussion and defence.

Inge. In this one case, there is no link between the things that could have been her reason for punishing and a fact that constitutes a sufficient reason to punish. In this case, Inge was innocent.⁵¹

In a reason-centric framework, each instance of just punishment would have to be a case in which there is a fact that gives the agent at least one reason to punish. The search for backup reasons or the desire to identify reasons with false propositions (a.k.a., alternative facts) is a manifestation of the mistaken belief that such pairings have to exist. They do not and that's why we should reject Normative Motivation. Once we drop that, we can call off the search. We can allow that an agent can do the right thing (e.g., punish an innocent child) even when there is nothing they could have had in mind that could have played the role of a sufficient or decisive normative reason to punish. We think that the right lesson to draw from three-option cases and our example is that evidence-relative reasons are not normative. Their presence and absence doesn't change what an agent ought to do once we have in place the relevant values and the agent's evidence.

When we reach a point where we can see no way to accommodate the normativity of rationality and of reasons, we think that we have to reject the Unification Thesis. The considerations that seem to show that the clash between reasons and rationality is inevitable even once reasons are epistemically constrained suggest to us that the agent ought to do the thing that's rational to do, not the thing that our best theory of normative reasons recommends. What would the alternative be? We don't see that there's anything wrong with this humble thought. There are principles of rationality. They identify ways of trying to cope with uncertainty, risk, and error in trying to see to it that our actions promote and respect various values and/or grounds of obligation. If they represent our best ideas about how to cope with uncertainty, risk, and error, we ought to do what they prescribe. We can see in advance that this means that we often ought to do what there is decisive evidence-relative reason not to do, but that is unfortunate. What is unfortunate, everyone seems to agree, is not always wrong. In this case, a principle of rationality that tells us to maximise expected value told us what to do, what to believe, and vindicated some naïve ideas about how some actions and attitudes ought to go together. Getting the details right required us to get clear on the values that matter to belief and to a certain kind of action, but we think it's a good thing that we can get a handle on our tricky case by using a familiar principle from decision-theory and some not wholly implausible claims about the values that matter in epistemic assessment and in our practice of blame and punishment.

⁵¹ Could Agnes* have sufficient reason (mass) to punish Inge? Yes, we might say that. But we had better not say that there was a reason (count) to punish, one that constituted an evidence-relative reason for Agnes* to punish Inge. (See §4.)

If our proposal is right, it vindicates the idea that the person who punishes and blames is right to do so when it's rational for them to have the beliefs that would make it seem that there were entries on the ledger under the column of things that counted in favour of punishing. Their existence, however, isn't essential to our explanation as to why these agents acted rightly. Agnes* acts rightly when she punishes 666 and when she punishes Inge. She will *think* that the ledger includes their guilt in both cases, but we know better. The right explanation as to why Agnes* ought to punish will appeal to things that might not be known to her, but things known to us (e.g., things about the relevant values, things about probability, things about how probabilities and values can work together to determine what she ought to do). Once we see this, we think that there's an interesting theoretical option to consider.

Because so many reasonologists have bought into the idea that reasons are facts that count in favour and look like familiar pros and cons, there might be some resistance to this idea, but we think it should be floated anyway. Perhaps we should think of normative reasons as things that figure in normative explanations. In our case, the explanations might be things like this. Agnes* ought to punish this child. Not because she's guilty, mind you. This child is innocent Inge. Agnes* ought to punish this child because doing so maximises expected value where the values that matter aren't just seeing to it that bad things happen to the guilty but that the guilty are punished by those who can rationally point to their guilt as the reason for which they are being punished. We wouldn't expect to find such considerations operative in the mind of Agnes. We find such considerations operative in the minds of theorists thinking about Agnes' situation from the outside and thinking about how people like Agnes ought to act in situations like this.

If we think about normative reasons this way, it seems evident that it would be a mistake to impose agent-relative epistemic constraints on such things. The things that would figure in our best theoretical accounts of why agents ought to do the things they ought to do might involve principles that no normal agent would ever think of and concepts that many perfectly virtuous agents might never grasp. Once we see that, we would think that Normative Motivation would have to be false when taken to be about normative reasons understood as things that figure in normative explanations. (Of course, on this conception of what a normative reason is, a normative reason needn't be a pro or a con, so it's important to remember that we're breaking from the conception of a reason that would be listed by a normal agent in a ledger when listing the pros and cons.) On this conception of normative reasons as facts that explain, appeals to platitudes about guidance would be misplaced. We don't think that the job of things that figure in such explanations is to provide an ordinary person with helpful guidance.

But why on earth did we think that normative reasons have to provide guidance? Normative reasons understood as facts that constitute pros and cons are, on the evidence-relative model, potential or actual objects of knowledge. And why would lists of such things be a useful guide? Normal agents don't distribute their probabilities uniformly across the objects of such a list. Decision-theory tells us that normal agents would be irrational to treat such lists of facts as guides for deciding what to do. If guidance matters, reasons understood as evidence-relative facts don't provide much by way of guidance. Rational degrees of belief, on the other hand, might be better suited for that. But there is no interesting way to recover a list of facts that constitute evidence-relative reasons from a description of the agent's rational degrees of belief.

References

- Alvarez, Maria. 2010. *Kinds of Reasons*. Oxford University Press.
- Alvarez, Maria. Forthcoming. Reasons for Action, Acting for Reasons, and Rationality. *Synthese*.
- Anderson, Charity and Hawthorne, John. Forthcoming. Knowledge, Practical Adequacy, and Stakes . In J. Hawthorne, & T. Gendler (Eds.), *Oxford Studies in Epistemology*, Volume 6. Oxford University Press.
- Blome-Tillmann, Michael. 2017. 'More Likely Than Not' - Knowledge First and the Role of Statistical Evidence in Courts of Law. In A. Carter, E. Gordon, and B. Jarvis (ed.), *Knowledge First: Approaches in Epistemology and Mind*. Oxford: Oxford University Press, pp. 278-292.
- Broome, John. 2005. Does Rationality Give Us Reasons? *Philosophical Issues* 15: 321-37.
- Broome, John. 2007. Does Rationality Consist in Responding Correctly to Reasons? *Journal of Moral Philosophy* 4: 349-74.
- Broome, John. 2013. *Rationality Through Reasoning*. Blackwell-Wiley.
- Buchak, Laura. 2014. Belief, Credence, and Norms. *Philosophical Studies* 169: 1-27.
- Christensen, David. 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: Oxford University Press.
- Cohen, Stewart. 1984. Justification and Truth. *Philosophical Studies* 46: 279-95.
- Comesaña, Juan. Forthcoming. A Plea for Falsehoods. *Philosophy and Phenomenological Research*.
- Comesaña, Juan and McGrath, Matthew. 2014. Having False Reasons. In C. Littlejohn and J. Turri (ed.), *Epistemic Norms: New Essays on Action, Belief, and Assertion*. Oxford University Press, pp. 59-81.
- Dancy, Jonathan. 2000. *Practical Reality*. Oxford University Press.
- Dietz, Christina. Forthcoming. Reasons and Factive Emotions. *Philosophical Studies*.
- Dorst, Kevin. Forthcoming. Lockeans maximize expected accuracy. *Mind*.
- Drake, Jonathan. 2018. Motivating Reason to Slow the Factive Turn in Epistemology. In V. Mitova (ed.), *The Factive Turn in Epistemology*. Cambridge University Press, pp. 177-92.
- Dutant, Julien. Forthcoming. Knowledge-Based Decision Theory and the New Evil Demon. In J. Dutant and F. Dorsch (ed.), *The New Evil Demon*. Oxford University Press.
- Dutant, Julien and Fitelson, Branden. MS. Knowledge Centered Epistemic Utility Theory.
- Easwaran, Kenny. 2016. Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities. *Nous* 50: 816–33.
- Easwaran, Kenny and Fitelson, Branden. 2015. Accuracy, Coherence, and Evidence. In T. Gendler and J. Hawthorne (ed.), *Oxford Studies in Epistemology*, Vol. 5. Oxford University Press: 61-96.
- Fantl, Jeremy and McGrath, Matthew. 2009. *Knowledge in an Uncertain World*. Oxford University Press.

- Fogal, Daniel. 2016. Reasons, Reason, and Context. In E. Lord and B. Maguire (ed.), *Weighing Reasons*. Oxford University Press, pp. 74-104.
- Foley, Richard. 2009. Belief, Degrees of Belief, and the Lockean Thesis. In F. Huber and C. Schmidt-Petri (ed.), *Degrees of Belief*. Springer, pp. 37–47.
- Gardiner, Georgi. 2018. Evidentialism and Moral Encroachment. In K. McCain (ed.) *Believing in Accordance with the Evidence*. Springer, 169-99.
- Goldberg, Sandy. 2017. Should have Known. *Synthese* 194: :2863-2894.
- Goldman, Alvin. 1986. *Epistemology and Cognition*. Harvard University Press.
- Goldman, Alvin. 2009. Williamson on Knowledge and Evidence. In P. Greenough and D. Pritchard (ed.), *Williamson on Knowledge*. Oxford University Press, 73-92.
- Graham, Peter A. 2010. In Defense of Objectivism about Moral Obligation. *Ethics* 121: 88–115.
- Gregory, Alex. 2016. Normative Reasons as Good Bases. *Philosophical Studies* 173: 2291–2310.
- Harman, Gilbert. 1968. Knowledge, Inference, and Explanation. *American Philosophical Quarterly* 5: 164-73.
- Ichikawa, Jonathan Jenkins. 2018. Internalism, Factivity, and Sufficient Reason. In V. Mitova (ed.), *The Factive Turn in Epistemology*. Cambridge University Press, pp. 66-83.
- Jackson, Frank. 1991. Decision-Theoretic Consequentialism and the Nearest and Dearest Objection. *Ethics* 101: 461-82.
- Jackson, Frank and Smith, Michael. 2016. The Implementation Problem for Deontology. In E. Lord and B. Maguire (ed.), *Weighing Reasons*. Oxford University Press, pp. 279-91.
- Kiesewetter, Benjamin. 2016. You Ought to ϕ Only if You May Believe that You Ought to ϕ . *Philosophical Quarterly* 66: 760-782.
- Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. Oxford University Press.
- Kiesewetter, Benjamin. 2018. How Reasons are Sensitive to Available Evidence. In C. McHugh, J. Way, and D. Whiting (ed.), *Normativity: Epistemic and Practical*. Oxford University Press: 90-115.
- Kiesewetter, Benjamin. Forthcoming. Review of *The Value of Rationality* by Ralph Wedgwood. *Ethics*.
- Lasonen-Aarnion, Maria. 2010. Unreasonable Knowledge. *Philosophical Perspectives* 24: 1-21.
- Leitgeb, Hannes. 2017. *The Stability of Belief: How Rational Belief Coheres with Probability*. Oxford University Press.
- Lempert, Richard. 1977. Modeling Relevance. *Michigan Law Review* 75: 1021–57.
- Littlejohn, Clayton 2012. *Justification and the Truth-Connection*. Cambridge University Press.
- Littlejohn, Clayton. Forthcoming. Being More Realistic About Reasons. *Philosophy and Phenomenological Research*.
- Littlejohn, Clayton. Forthcoming b. Truth, Knowledge, and the Standard of Proof in Criminal Law. *Synthese*.
- Lord, Errol. 2018. *The Importance of Being Rational*. Oxford University Press.
- Makinson, D.C. 1965. The Paradox of the Preface. *Analysis* 25: 205-7.
- Mantel, Susanne. 2017. Worldly Reasons: An Ontological Inquiry into Motivating Considerations and Normative Reasons. *Pacific Philosophical Quarterly* 98: 5-28.
- Moeller, Emil. 2015. *Defeasible Knowledge*. Oxford University DPhil Thesis.
- Moss, Sarah. Forthcoming. *Probabilistic Knowledge*. Oxford University Press
- Nelkin, Dana. 2000. The Lottery Paradox, Knowledge, and Rationality. *The Philosophical Review* 109: 373-409.
- Neta, Ram. 2009. Treating Something as a Reason for Action. *Nous* 43: 684-99.
- Olsen, Kristian. 2017. A Defense of the Objective/Subjective Moral Ought Distinction. *The Journal of Ethics* 21: 351–373.
- Owens, David. 2000. *Reason without Freedom*. Routledge.
- Parfit, Derek. 2011. *On What Matters*. Oxford University Press.
- Pollock, John. 1986. The Paradox of the Preface. *Philosophy of Science* 53: 246-58.

- Prichard, H.A. 2002. Duty and Ignorance of Fact. In J. MacAdam (ed.) *Moral Writings*. Oxford University Press, pp. 84-101.
- Redmayne, Michael. 2008. Exploring the Proof Paradoxes. *Legal Theory* 14, 281–309.
- Regan, Don. 1980. *Utilitarianism and Cooperation*. Oxford University Press.
- Ross, W.D. 1930. *The Right and the Good*. Oxford University Press.
- Ryan, Sharon 1991. The Preface Paradox. *Philosophical Studies* 64: 293-307.
- Ryan, Sharon. 1996. The Epistemic Virtues of Consistency. *Synthese* 109: 121–41.
- Schroeder, Mark. 2008. Having Reasons. *Philosophical Studies* 139: 57-71.
- Smith, Holly. 2010. Subjective Rightness. *Social Philosophy and Policy* 27: 64-110.
- Smith, Martin. 2016. *Between Probability and Certainty: What Justifies Belief*. Oxford University Press.
- _____. Forthcoming. When Does Evidence Suffice for Conviction? *Mind*.
- Sosa, Ernest. 2015. *Judgment and Agency*. Oxford University Press.
- Steele, K. MS. Crime, Punishment, and 'Individualized' Evidence.
- Sutton, Jonathan. 2007. *Without Justification*. MIT University Press.
- Sylvan, Kurt and Sosa, Ernest. 2018. The Place of Reasons in Epistemology. In D. Star (ed.), *The Oxford Handbook of Reasons and Normativity*, pp. 555-75.
- Thomson, Judith. 1990. *The Realm of Rights*. Harvard University Press.
- Titelbaum, Michael. Forthcoming. Reason without Reasons For. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Vol. 13. Oxford University Press.
- Unger, Peter. 1975. *Ignorance*. Oxford University Press.
- Way, Jonathan and Whiting, Daniel. 2016. Reasons and Guidance (Or, Surprise Parties and Ice Cream). *Analytic Philosophy* 57: 214-35.
- Way, Jonathan. 2017. Reasons as Premises of Good Reasoning. *Pacific Philosophical Quarterly* 98: 251-70.
- Wedgwood, Ralph. 2002. Internalism Explained. *Philosophy and Phenomenological Research* 65: 349-69.
- Wedgwood, Ralph. 2016. Objective and Subjective 'Ought'. In N. Charlow and Matthew Chrisman (ed.), *Deontic Modality*. Oxford University Press, pp. 142–68.
- Wedgwood, Ralph. 2017. *The Value of Rationality*. Oxford University Press.
- Whiting, Daniel. 2014. Keeping Things in Perspective: Reasons, Rationality, and the Apriori. *Journal of Ethics and Social Philosophy* 8: 1-23.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford University Press.
- Williamson, Timothy. 2009. Replies to Critics. In P. Greenough and D. Pritchard (ed.), *Williamson on Knowledge*. Oxford University Press, 279-85.
- Williamson, Timothy. 2017. Very Improbable Knowing. *Erkenntnis* 79: 971-99.
- Worsnip, Alex. 2016. Belief, Credence, and the Preface Paradox. *Australasian Journal of Philosophy* 94: 549-62.
- Zimmerman, Michael. 2008. *Living with Uncertainty*. Cambridge University Press.