

Do Reasons and Evidence Share the Same Residence?

Clayton Littlejohn
King's College London

Introduction

It is striking fact that so many good people know to do the things that should be done even when they know very little about what ultimately makes these things the things to do.¹ They know to call their parents (but not something insulting) and to feed their children (but not to tigers) without the benefit of the theoretical understanding they'd pick up in an ideal ethical theory class. How could this be?

Star suggests that this is largely down to two things. First, these good people are the good people because they are appropriately responsive to normative reasons (2015: 97). These agents are moved by good reasons, albeit derivative reasons that are distinct from the fundamental reasons that would figure in the right theory of the right. These derivative reasons reliably lead these agents to do the right thing because their status as reasons ensures that they are *evidence* that these agents ought to perform the relevant acts. Thus, while the good folk might not act for the reasons that figure in good theories of right action, it isn't an accident that they act as these theories tell us they ought to.

Star's explanation of our striking fact appeals to the reasons as evidence thesis:

R: Necessarily, p is a reason for A to ϕ if and only if the fact that p is evidence that A ought to ϕ (2015: 13).²

The thesis tells us that facts that constitute normative reasons do so by being evidence that the agent ought to do the things that they constitute reasons to do. The main rival to Star's approach takes reasons to be reasons by virtue of standing in *explanatory* relations to normative facts rather than some sort of evidential relation. An important virtue of Star's approach is that this rival approach struggles to explain our striking fact.

Our discussion focuses on R. It's natural for reasonologists considering R for the first time to think that its defenders are confused. Reasons, they'll say, are the kinds of things that *make* right things right and *determine* what we ought to do (and believe and feel). They don't serve as *indicators* or *clues* that we ought to do these things. If they were clues or indicators, there would have to be some further independently operating things that made things right and *these* further things would seem to play the theoretical role of normative reasons. The clues and indicators are too far removed from the action to be normative reasons. This worry is illustrated by thinking about some (possibly fallible, possibly infallible) guru telling you that you ought to ϕ . While that might be *evidence* that you ought to ϕ , it surely cannot be a

¹ I would like to thank Maria Alvarez, Christina Dietz, Julien Dutant, John Hawthorne, David Owens, and Daniel Star for the countless hours spent discussing reasons. Julien and Daniel were particularly helpful in working through the details of this paper.

² R was introduced into the literature by Kearns and Star (2008, 2009). Readers familiar with these papers can learn a great deal from Star (2015) as the monograph breaks new ground and doesn't reuse much of the previous material.

good reason to ϕ , could it? At best, doesn't it indicate that something else counts as a good reason to ϕ and, if the advice is good, make it the case that you ought to ϕ ?³

Some of what I say *might* be in the spirit of this objection, but this objection isn't convincing as stated. First, there's the *non sequitur*. There's nothing in R that precludes the things that constitute evidence being among the things that make the right things right. Reasons might wear two hats. It's possible that the fundamental reasons are fundamental because they wear both of them. Second, one might worry that this objection overgeneralizes. Consider the indirect consequentialist. The indirect consequentialist would probably want to allow that there are genuine reasons to ϕ that are distinct from but derived from the fundamental reasons that explain why agents ought to act in certain ways. You might think there's a sense in which such derivative reasons wouldn't be the kinds of things that figure in fundamental explanations of the ought-facts, the kinds of things that *make* right acts right, so straight off we should see that the test to determine whether a reason is a genuine normative reason or not probably isn't to ask whether it figures in fundamental explanations. And once that point is acknowledged, the indirect consequentialist might worry that they won't have any good story to tell about how such reasons could be both genuine normative reasons that are distinct from the fundamental ones and how it could be that these are the reasons that motivate the good people. They might be relieved in reading through *Knowing Better* that Star has plausible way of making sense of all this.

There might still be a problem with R that's related to this concern. It might be something that proponents of the main rival view, the view that treats reasons as things that figure in explanations of oughts, could appeal to. I don't think I've found the fatal flaw in anything central to Star's system, but I do think that some serious work will have to be done to dodge difficult choices.

Reasons, Evidence, and Epistemic Constraints

Star thinks that normative reasons and pieces of evidence consist of facts. R tells us the relation they have to bear to prospective actions to constitute reasons for them. Upon first encountering this idea that reasons to ϕ do so by virtue of being evidence that the agent ought to ϕ , I thought Kearns and Star were putting forward the view that that reasons that bear on whether some agent should ϕ consisted of the facts that this agent possessed as part of her evidence (e.g., the facts that she knew), but Star was quick to correct me on this. R allows for the possibility that the facts that, say, gave Agnes a reason to ϕ were all pieces of evidence that Agnes ought to ϕ even if they weren't part of *her* evidence. Indeed, they might be reasons without being part of *any* subject's evidence. Star has options. If you think there can be reasons that bear on whether to ϕ that nobody knows or could know, R doesn't rule this out. R only says that such things must, *inter alia*, be evidence. If, however, you think that there can only be reasons that belong to some subject's evidence (e.g., the agent's evidence), R doesn't rule this out.

R does seem to rule some options out. It seems to rule out views that violate this thesis:

³ See Broome (2008) as well as McNaughton and Rawling (2011). See Kearns and Star (2011) for responses.

Commonality Thesis: The set of things that stand in some specific reason-relation (e.g., the things that constitute reasons for A to ϕ) are all things that stand in some evidential-support relation (e.g., things that constitute evidence that A ought to ϕ).⁴

If normative reasons are pieces of evidence, reasons and evidence share the same residence. It might be that such things aren't subject to any epistemic constraints at all so that there's no limit on the kinds of facts that might constitute reasons for Agnes to ϕ (e.g., they might include the known to be known, known, merely knowable, unknowable, unknowably unknowable, etc.). It might be that the only reasons that matter are things that belong to some agent's evidence (e.g., it might be the only facts that stand in the reason-relation and count as reasons for Agnes to ϕ are all known by her). On this view, all normative reasons are *possessed reasons* or reasons that some subject *has*, so that Star's R collapses into:

R* Necessarily, p is a reason A has to ϕ if and only if p is evidence that A has that A ought to ϕ (2015: 13).⁵

On this view, all the reasons that stand in the relevant reason-relation are possessed as pieces of evidence by the agent in question:

Epistemic Containment: A fact is a normative reason only if it is a fact that belongs to some body of evidence by virtue of some subject bearing an interesting epistemic relation to it (e.g., knowledge, possible knowledge, etc.).⁶

While Star doesn't seem to think that R collapses into R* and doesn't defend Epistemic Containment, we'll see that there's a *prima facie* plausible argument from R to Epistemic Containment. While Epistemic Containment isn't a wildly unpopular view, we'll also see that it's not wholly unproblematic. It seems to me that the best way to avoid these difficulties might be to reject the Commonality Thesis. While it's easy to see why Star's rivals might reject that, it's harder to see how, given R, the Commonality Thesis might be false.

From R to Epistemic Containment

One of R's selling points is that it's supposed to help us see how it's possible to weigh reasons against each other in the course of practical deliberation. Doing so is simply a

⁴ Notice that the Commonality Thesis is compatible with the view that *epistemic reasons* are subject to epistemic constraints that other reasons are not. Skorupski defends such a view, but I've (2012) argued that this view is mistaken in defending a version of the Enkratic Requirement.

⁵ When someone says that Agnes has a reason to ϕ , I understand the speaker to be saying that there's a reason for Agnes ϕ (no more, no less) *and* I don't hear *that* as implying much of anything about Agnes' epistemic access to the relevant reason. In this discussion, we're assuming that the reasons that an agent has to ϕ are all reasons there are for this agent to ϕ *and* that this agent bears an interesting epistemic relation to this reason, the kind of relation that the agent would bear to something if that thing were part of her evidence. Like Star, I take this relation to be knowledge.

⁶ For helpful discussions of Epistemic Containment and motivations for it, see Gibbons (2013), Mason (2013), and Zimmerman (2008).

matter of weighing evidence, the evidence that we ought to ϕ against evidence to the contrary. Since the process of weighing evidence is non-monotonic, R seems to explain the non-monotonicity of practical deliberation. If R explains this, it should also vindicate Ross' observation that it's possible to judge that we should ϕ all things considered and still rationally regret that we had to ϕ . We can regret that we had to ϕ because we can still see that there were good reasons to do otherwise.⁷

While R seems ideally suited to vindicate Ross' observation, it only seems to do so if combined with Epistemic Containment. To see why, suppose that it's a fact that A ought to ϕ . If so, this fact should be part of the evidence there is. If it is evidence for anything, it should be evidence for itself. (It's natural to think that $P(A \text{ ought to } \phi | A \text{ ought to } \phi) > P(A \text{ ought to } \phi)$, in which case the fact that A ought to ϕ is evidence that A ought to ϕ .⁸) If it *is* evidence for itself, it follows from the fact that A ought to ϕ that $P(A \text{ ought to } \phi) = 1$. (Remember that we're assuming that P is defined over the evidence there is.) If so, the evidence there is can include *nothing* that is evidence that A ought not ϕ or ought to do other than ϕ because $P(A \text{ ought not } \phi) = 0$, in which case there's nothing in the evidence that's evidence for this.⁹

It might have seemed that R would provide a neat explanation of Ross' observations about rational regret. The process of weighing the evidence is non-monotonic. Unfortunately, the *reflexivity* of evidential support causes difficulty on this front. Once the fact that A ought to ϕ is included in the evidence, it's no longer possible to think of any consideration as evidence that counts against this. Thus, if R is going to help us understand how a rational agent can feel the force of the case for ϕ -ing *and* the case against it, we'll either have to distinguish reasons from evidence or restrict the things that get into the evidential support relation to allow that there might be both evidence that someone ought to ϕ and evidence that this person ought not ϕ .¹⁰

⁷ See Ross (1930). For a helpful discussion of this and the costs and benefits of thinking of the support of reasons in terms of evidential support, see Hawthorne and Magidor (forthcoming.).

⁸ If 'the evidence' includes all and only the facts, it's quite possible that the inequality doesn't hold. It might well be that A ought to ϕ is true only if there's some *other* fact or set of facts that entails that A ought to ϕ (e.g., the fact that A ought to ϕ and the moon isn't made of cheese). If so, it would still follow that A ought to ϕ only if $P(A \text{ ought to } \phi) = 1$.

⁹ If there were evidence for some hypothesis h when $P(h) = 0$, there would need to be something in the evidence such that $P(h|e) > P(h)$. There could be no such evidence to generate this inequality. This is because $P(h|e)$ is defined as $P(h \& e)/P(e)$ and $P(h \& e) = 0$ when $P(h) = 0$.

¹⁰ Might Star avoid these difficulties by saying that the evidential support relation is not defined over something like a subject's evidence rather than the total evidence? Perhaps, but then the resulting view would not seem to differ in any interesting way from the views that identify the evidence with either the knowable facts or the known facts. I think he'd want to avoid views on which something can stand in the reason-relation and do the work done by normative reasons (e.g., make it right for A to ϕ , make it the case that A ought to ϕ , etc.) without being evidence.

Isn't there an obvious way to fix this? Even if, say, $P(p) = 1$ given *some* body of evidence, there might be some *other* body of evidence on which $P(p)$ is less than 1. This is all we need to make a process like weighing the evidence/weighing the reasons non-monotonic. Provided that Star's view allows for different bodies of evidence and thus different support relations between, say, p , and two different bodies of evidence, can't we understand how it's possible *both* that some agent ought to ϕ and this agent is able to weigh evidence for/against the hypothesis that she ought to ϕ and reasons that count for/against ϕ -ing?¹¹

This would fix the problem but it raises a tricky issue. Shouldn't there be some privileged set of *reasons* that either determines that Agnes ought to ϕ or plays a unique role in the supervenience base of the fact that she ought to ϕ ? We're comfortable with the idea that there might be different bodies of evidence, that there's a perfectly good notion of probability according to which p is the probability of p waxes and wanes as the set of propositions in some body of evidence changes over time, and with the idea that there's not some privileged body of *evidence* that gives us p 's real probability (whatever that might be). Should we be comfortable saying similar things about normative reasons? I don't think we should be comfortable with the idea that there are diverse sets of reasons where, say, relative to one set Agnes ought to ϕ when there's some other set where relative to those it is not the case that Agnes ought to ϕ . There should be a privileged set of normative reasons.

The natural way to make sense of different bodies of evidence is to think about the individuals, the facts, and the epistemic relations that hold between them. The epistemic relation (e.g., knowledge) would function as a kind of restrictor that rules out some facts from figuring in the set of evidence so that we don't end up saying that there's some set of evidence in which all truths have probability 1 *and* help us identify the interesting sets of facts. Perhaps, then, we should try something similar and let the agent's epistemic relations identify some privileged set of facts to be normative reasons that count as the reasons that apply to her.

If we impose restrictions on the sets of facts that constitute some body of evidence, it seems to follow, given R, that we're going to impose the same constraints on the set of facts that constitute some body of normative reasons. This is what the Commonality Thesis says. This leads us to Epistemic Containment. Thus, to understand how a process like practical reasoning could be non-monotonic, proponents of R should say that evidential support relations aren't defined over the totality of facts (or 'the evidence there is') but *only* over evidence that's either possessed or potentially possessed. Moreover, it must be evidence that's possessed *by* the relevant agent, not some omniscient observer. This, in turn, constrains the things that can stand in the reason-relation. The reasons that stand in specific reason-relation (e.g., they are reasons for Agnes to ϕ) have to be possessed evidence and assuming there's some privileged set it's natural to think that these reasons are part of Agnes' body of evidence.

¹¹ Applied to our case, the idea would be that while it follows from the fact that A ought to ϕ , $P_{\text{the evidence there is}} = 1$, it wouldn't follow from this fact that $P_{\text{Agnes' evidence}} = 1$, so Agnes could still weigh evidence for and against, provided that Agnes' evidence doesn't include everything in the evidence there is (and specifically does not include the fact that she ought to ϕ .)

Against Epistemic Containment

While I think *evidence* has to be epistemically contained, I don't think the same is true of *normative reasons*. I'll briefly discuss two problems with Epistemic Containment.

First, Star and I both like the idea that a person shouldn't treat p as a reason for the purposes of practical deliberation if she doesn't know p . If we suppose that this is correct, suppose that 'ought' implies 'reason', and suppose that it's possible for someone to treat p as if it's a reason for action when p is an unknowably unknowable proposition, we have to reject Epistemic Containment. There would be a reason for this subject not to treat p as a reason, but this reason wouldn't be a fact the subject possessed as evidence. A subject's evidence just is her knowledge.

Second, we need an account of *how* the reasons that count as reasons for Agnes to do things relate to the facts about what Agnes ought to do. Here are some natural thoughts. First, if there's a strong reason for Agnes to ϕ and no reason for her not to ϕ , she ought to ϕ . Second, when there are competing reasons that pull her in different directions, she ought to act in accordance with the stronger reason. In a framework that includes Epistemic Containment, these apparent platitudes about reasons generate strange results.

Consider:

Lights: Agnes is responsible for managing the lights in a large building. She has very strong evidence that she ought to illuminate room 1. She has very strong evidence that she ought to illuminate room 2. The situation is the same for room 3, room 4, room 5, ..., and (finally) room 1,000,000. She can illuminate all the rooms at once. She can do it one by one. Just as she's about to illuminate some rooms when her guru tells her that she ought not illuminate all the rooms because there is precisely one room that she ought not illuminate. As often happens in this kind of situation, he drops dead before he can explain why this is and identify the room. Now Agnes has very strong evidence that she ought not illuminate all the rooms *and* also has very strong evidence for each room that she ought to illuminate it.

Assuming that the strength of the reasons there are determines what Agnes should do, this set of reasons is identified with Agnes' knowledge, and that their strength is understood in terms of strength of evidential support, we get the result that Agnes ought to ϕ if the degree of evidential support for the proposition that she ought to ϕ is sufficiently high. When we think about preface-like situations, we can imagine that Agnes' evidence provides sufficiently strong support for (a) believing of each room that it ought to be illuminated *and* (b) that she ought not illuminate all the rooms. So, it seems that for each room Agnes ought to illuminate it *and* it is not the case that she ought to illuminate all the rooms. If 'ought' agglomerates, we get a dilemma. (Because of (a), she ought to illuminate all the rooms if she ought to illuminate each room.) And if 'ought not ϕ ' implies 'not ought ϕ ' we get a contradiction.

If we want to avoid this kind of problem, we can retain the idea that the strength of reasons determines what a subject should do *provided* that we find something outside the set of things Agnes knows that tells us whether she really

ought to illuminate every room or which room(s), if any, should not be illuminated. This, in turn, forces us to deny Epistemic Containment.

Discussion

The general shape of the argument against R is this:

1. All evidence is possessed evidence.
2. If so, the Commonality Thesis implies that all the normative reasons that are reasons for A to ϕ are reasons that A has.
3. Not all the normative reasons for A to ϕ are reasons that A has.
4. Thus, the Commonality Thesis is mistaken.
5. If the Commonality Thesis is mistaken, so is R.
6. Thus, R is mistaken.

I don't take this argument to be decisive, by any means. At best, it points to some potential trade-offs we might have to make if we accept R. There are a number of things Star might say to try to resolve the difficulties raised above.

To understand how there could be norms like the knowledge norm of practical reason and understand how there could be misleading evidence about what an agent ought to do, we'd want to allow that the set of things that stand in the reason-relation to include things that couldn't be included in a subject's evidence. Theoretically, if we decided to retain this kind of norm we'd have to do one of three things. We might simply abandon R and look to an alternative such as the idea that reasons are things that figure in explanations. This is the least attractive option for Star. Second, we might retain R and reject Epistemic Containment. We could allow that just about any old fact might potentially constitute a normative reason that counts as a reason for some specific subject to ϕ and accept that this means that the evidential support relation is defined over things that are not and could not be included in this subject's evidence. If we go this route, we would have to try to make sense of the idea of unpossessed evidence. Finally, we might retain R, retain some version of Epistemic Containment, and try to make sense of a view on which the reason-relation is defined over a set of facts that do not and could not stand in the evidential support-relation. This would allow the strongest reason to determine what a subject ought to do, but it seems to come at the cost of allowing reasons to be things that could not potentially enter into evidential support-relations. It is difficult to see how this theoretical option could be open for a proponent of R, but I don't have a proof in my pocket that shows that this could not be done.

Because rival accounts of reasons don't take reasons to be evidence, they don't have much reason to accept the Commonality Thesis. One potential advantage of the view that identifies reasons with things that enter into explanations is precisely that they have the resources to handle cases like Lights. There's something in the universe that determines what Agnes ought to do, we might never know what it is, this reason explains why some of Agnes' evidence is misleading, and there's no reason to think that these reasons have to be epistemically accessible in the way that we take evidence to be. They are inaccessible but potent.

If we set the case aside, there's an interesting and important issue that potentially divides the reasons as evidence view from the explanatory view of reasons. It strikes me that the best direction for Star to go would be to embrace Epistemic

Containment and embrace the kind of view that doesn't allow that there are facts that stand in the reason-relation that aren't accessible to the agent. In so doing, he could just deny that there are norms like the knowledge norm that imply there are potent, potentially decisive reasons that aren't things we could ever hope to know. Star might find theoretical support for his view by thinking about whether reasons do their work by being the kinds of things that can guide us by figuring in deliberation.¹²

I've always been more sympathetic to the view that says that normative reasons are tools that we use in our normative theories and that it didn't matter whether the considerations that figured in our theories were accessible to agents or could help agents see what they ought to do. There's another school of thought, though, that sees reasons as things that are there to guide agents by playing a role in reasoning, and it might be that this approach is both a better fit for R and the kind of thing that, if properly defended, would help us see the advantages of the reasons as evidence thesis over its rivals.

References

- Alvarez, Maria. 2010. *Kinds of Reasons*. Oxford University Press.
- Broome, John. 2008. Reply to Southwood, Kearns and Star, and Cullity. *Ethics* 119: 96-108.
- Gibbons, John. 2013. *The Norm of Belief*. Oxford University Press.
- Hawthorne, John and Ofra Magidor. Forthcoming. Reflections on Reasons. In D. Star (ed.), *Oxford Handbook of Reasons and Normativity*. Oxford University Press.
- Kearns, Stephen and Daniel Star. 2008. Reasons: Explanations or Evidence? *Ethics* 119: 31-56.
- Kearns, Stephen and Daniel Star. 2009. Reasons as Evidence. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 4, pp. 215-42.
- Kearns, Stephen and Daniel Star. 2011. On Good Advice: a Reply to McNaughton and Rawling. *Analysis* 71: 506-8.
- Littlejohn, Clayton. 2012. *Justification and the Truth-Connection*. Cambridge University Press.
- Littlejohn, Clayton. 2013. The Russellian Retreat. *Proceedings of the Aristotelian Society* 113: 293-320.
- McNaughton, David and Piers Rawling. 2011. The Making/Evidential Reason Distinction. *Analysis* 71: 100-02.
- Mason, Elinor. 2013. Objectivism and Prospectivism about Rightness. *Journal of Ethics and Social Philosophy*. 7: 1-21.
- Ross, W.D. 1930. *The Right and the Good*. Oxford University Press.
- Skorupski, John. 2010. *The Domain of Reasons*. Oxford University Press.
- Star, Daniel. 2015. *Knowing Better: Virtue, Deliberation, and Normative Ethics*. Oxford University Press.
- Way, Jonathan and Daniel Whiting. Forthcoming. Reasons and Guidance (Or, Surprise Parties and Ice-Cream). *Analytic Philosophy*.
- Zimmerman, Michael. 2008. *Living with Uncertainty*. Cambridge University Press.

¹² For mostly sympathetic discussions of the guidance idea, see Gibbons (2013), Skorupski (2010), and Way and Whiting (forthcoming). For less sympathetic discussion, see Littlejohn (2013).