

Normativity and Contrastive Explanation

David Olbrich

Submitted for admission to the degree of
Master of Philosophical Studies

2015

University College London



Declaration

I, David Daniel Olbrich, confirm that the work presented in this thesis is my own. I have properly acknowledged all ideas drawn directly from other authors.

Wednesday, 26th August, 2015

Abstract

My thesis concentrates on the distinction between *pro tanto* reasons and all-things-considered judgments, and their relation to normative justification. Negatively, it seeks to show that a prevailing kind of account of this relation should be rejected, namely that family of views which takes it that every reason has an associated weight, and the truth with respect to any issue is established by which set of reasons is weightiest. Through an examination of Ross' doctrine of *prima facie* duties, this discussion also leads to a formulation of the central problem which any account of this relation must seek to solve. Positively, this thesis develops a new account of the relation between *pro tanto* reasons and all-things-considered judgements, based on the fundamental insight that a justification of normative propositions is identical to an explanation of their truth, were they to be true. I defend this identity claim, and seek to generate an account of justification from an account of explanation.

Drawing on a deservedly popular 'contrastive' conception of explanation in the philosophy of science, I show how we can fruitfully think of justification as itself contrastive. Part of this is showing how the notion of a burden of explanation can shed light on the notion of a burden of justification, so a conception of justification emerges according to which a justification for a normative proposition consists in an solution to all those burdens of justification which it incurs. In turn, this feeds a conception of reasons, and their role in justification, alternative to that envisaged in a weighing model: *pro tanto* reasons determine the correct all-things-considered judgment insofar as they determine to what extent the truth of that judgement has an adequate explanation, such that the correct all-things-considered judgement is just that judgement whose truth would have a fully adequate explanation.

Acknowledgments

This thesis would not have been possible without the many contributors to both its content and to my intellectual development over the last two years. I would especially like to thank my supervisor Douglas Lavin, whose tireless engagement with my work led me to recognise that there was far more in what I was saying than what I initially thought. This thesis also benefited from early supervision by Lucy O'Brien. I would also like to thank Edgar Phillips and Sean Whitton for their thoughts on my work. Above all I would like to acknowledge the wonderful environment of the UCL Philosophy Department graduate scene, which has never failed to stimulate or motivate me.

Contents

1. Reasons and Justification	6
2. The Weighing Model: Origins, Motivations and Problems	23
3. Burdens of Justification	42
4. Contrastive Explanation in the Structure of Justification	70
Bibliography	81

Reasons and Justification

We make normative claims: claims about what anyone of us should do, should feel, or what we should believe. How can anyone justify such claims? There are two questions here. What calls for justification? How can justification itself answer this call? A start to the second question is that justification requires reasons. We can justify normative claims only by showing that there are reasons that count in favour of their being true. For example, we can only justify the claim that it is wrong to prank the insecure by adducing the reasons for it, namely that doing so inflicts severe psychological torment, and is almost never funny. An improvement on this formulation is that we can justify normative claims only by showing that all the reasons, taken together, count in favour of their being true. So, even though by refraining from pranking the insecure one excludes them from more 'normal' elements of social interaction and so from a social in-group, still, taken together with the fact that doing so inflicts severe psychological torment, these reasons overall suggest that it is wrong to prank the insecure. What then does it take for the reasons to 'overall' count in favour of some claim? This opens up two lines of inquiry. The first investigates what reasons are and how they work. The second investigates how we need them to work, given what justification requires of them.

This chapter proceeds as follows. Initially I will demarcate more precisely what phenomenon I mean to focus on within what is ordinarily called 'justification'. I will then argue that justification of normative propositions (as opposed to other possible objects of justification) forms a distinctive kind in that justification of this kind is also explanation. On the basis of this equivalence I then set up the fundamental problem this thesis aims to confront, which concerns the structure of justification.

'Justification' is polysemous: one thing it means is the act or process of justification, something that occurs at a time and place, carried out by a justifier, to another or to themselves. 'Karl justified himself to his boss in the common room and was out in time for lunch' is an example of this use.

Throughout this thesis, however, I will be concerned with another use of ‘justification’, justification as the abstract entity which is argued over or thought about.¹ (The word ‘thought’ appears to work much the same way.) ‘Was there a justification for Karl’s actions?’, ‘Karl’s justification for his actions was that he was getting the firm the best deal he could’, ‘How are justifications structured?’, these are examples of ‘justification’ used this way. This is how I will use ‘justification’. Using both senses, we could say when we engage in justification, we aim to lay out an adequate justification in such a way that the person to whom the justification is addressed accepts it as adequate. It is important to distinguish this from a third use where ‘justification’ is a property of propositions, or beliefs, consisting in the degree to which they are justified, a sense of ‘justification’ which I shall not use.²

Another clarificatory point concerning my use of ‘justification’: even with justifications in the abstract sense, there is a way in which justification is relative to the epistemic state of the inquirer. If your evidence overwhelmingly points to Captain Scarlett’s being as destructible as normal humans, when in fact he is not, then we might say that relative to that evidential state, there is a justification for using a gun if one were to try to kill him. I do not doubt that this is a sensible use of ‘justification’, but still, when I talk of justification, I will always mean a justification detached from any relativity to evidential state, justification *tout court*. So, for instance, the fact that pranking the insecure inflicts severe psychological torment on them is the justification *simpliciter* for it’s being wrong to do it. This is not to say that justification cannot depend on facts about one’s evidential state, but merely that the justification itself does not take one’s evidential state as a parameter.

When we talk of the adequacy of a justification we are usually talking about justification in the abstract sense. A justification is adequate when it is supported by the reasons. With ‘Karl justified himself to his boss, giving

¹ Conceived in this way, justification is the kind of thing to which a non-trivially formulated Identity of Indiscernibles might apply: justifications are the same just in case they have the same intrinsic properties. If two justifications purport to justify the same thing with the same reasons, they are the same justification.

² Justification is used in this way in ‘Transmission of Justification and Warrant’, Moretti, Luca and Piazza, Tommaso, *The Stanford Encyclopedia of Philosophy*

adequate reasons', it does not follow that Karl adequately justified himself to his boss. Whether Karl adequately justified himself to his boss depends on his boss' level of uptake. If Karl gives his boss adequate reasons but his boss is a dolt and rejects the justification, then it will follow that there is an adequate justification for Karl's actions (and Karl had it) even though Karl could not adequately justify himself to his boss. Bound up with notions of the adequacy of a justification is the maximal use of 'justified', meaning 'fully justified'. To use 'justified' in this sense is to imply that the thing which is justified is fully justified. If something is adequately justified, then I use that to mean that it is fully justified. There is even a sense in which a justification fails to be a justification at all if it is not adequate: one might express scepticism about a purported justification with the phrase 'That's not really a justification.'

What counts as a justification depends on the kind of inquiry being pursued. Suppose two detectives are in a murder investigation, and one says 'The Colonel is the culprit.' The other detective asks for a justification. The first one replies, 'He had a motive, and his fingerprints were on the murder weapon, and no-one else had a motive, and no-one else could have had anything to do with the murder weapon.' This is a pretty good justification. It shows that the Colonel is almost certainly the culprit. A belief to that effect would be thereby vindicated; giving up that belief would be denied vindication. The relevant criterion here is truth indication. What made it a justification was that it indicated that it was true that the Colonel committed the crime. This seems platitudinous.

When justifying one's actions, in contrast, the relevant criterion for a justification is whether the action was right, legitimate or advisable. If Karl sells all the stock to a single bulk buyer at an inflated price, resulting in a high revenue but no potential repeat customers, for a justification of his actions to be successful it would have to show that it was a good move to make the deal, that the deal does not deserve to be regretted. Action descriptions are truth assessable: 'Karl got the firm the best deal he could' might be true or false. But to justify one's actions, as opposed to justifying a claim about what those actions were, is to show that they were right, legitimate or advisable (or possibly something else: I wish to impose no restriction). Again, this seems

platitudinous. This point applies more broadly: this criterion for a justification also holds when it is social arrangements being justified, or when justifying one's own existence.

What is the relevant criterion of a justification for normative propositions: propositions with a normative subject-matter? For example: 'It is good when people don't have to bottle up their emotions', or 'Perfidy should be absolutely forbidden.' This is (prima facie at least) truth-assessable so we might well think that the relevant norm of a good justification is truth indication, as with justification for non-normative propositions above. But 'it is right that people engage in poverty relief' is more closely linked to action, so we might think that, for this example, the norms of justification consist in showing the action of engaging in poverty relief to be right or legitimate. Similarly with propositions such as 'A high level of piety is best in a population.' The concern with social arrangements might suggest that the relevant criterion of a justification for such propositions is whether the arrangement referred to therein is right, legitimate or advisable. These would suggest that the criterion of a justification for normative propositions is, in general, for that which the proposition concerns to be shown to right, legitimate or advisable³. This would also be a natural, if not inevitable, consequence of meta-ethical views according to which the expression of a normative proposition can itself be understood as an action which bears no relation to a truth-assessable content.

³ Propositions concerning what we ought to believe are included as well, though their inclusion is slightly more problematic: an initially plausible criterion of adequacy for a justification for a proposition about what we ought to believe is what we should believe if it's true. In contrast, a good justification for a claim about how we ought to act doesn't consist in showing that the action is true—this is meaningless—so it looks as though we are forced to draw a distinction between the norms of justification for propositions concerning how we ought to act, and those for propositions concerning what we ought to believe. But it's not true that we ought to believe something iff it's true. All kinds of other considerations may intervene. We shouldn't believe every true proposition, for the sake of conserving memory or attention (e.g. see 'Clutter Avoidance', p.15, in Harman (1986)). Furthermore, there may be specific cases in which it's better to believe a falsity: say one is going into an important exam, it may be better for one's nerves and thereby one's performance if one believes that the exam is not very important. Whether one can knowingly believe falsities is a different question, but that does not refute the point that one ought to believe it if one can, say if one is not aware that it is in fact very important.

Normative propositions, then, appear to be stuck between these two platitudes. So a question emerges here over what is the correct construal of justification for normative propositions in general, if indeed there is an entirely general and unified kind of justification for all normative propositions.

What demarcates normative propositions is a certain criterion for what counts as a justification for them; but this justification must be more than truth indication as it was with the standard non-normative propositions. A starting indication for this lies in the case of using the moral testimony of another. Suppose someone believes that it is good when people spend their lives engaging in poverty relief, and we ask them for a justification of this claim; and they replied that the justification was that a wise person known to them says so. That person is reliable on the issue, so if they say so, it must be true. But this is an inadequate or inappropriate justification. Indeed it is not really a justification at all. There is an aspect in which it is a successful justification; *ex hypothesi* the wise person is reliable on this issue and the speaker is guided to the truth in taking up what they say. But that the wise person knows is not a justification for the truth of the claim; rather, if a justification for its truth exists at all, it must exist before it is possible for them to know and to communicate that knowledge. Hence the speaker has failed to give us a justification by mentioning the wise person's knowledge.

We may be tempted to think (in line with much of the literature on moral testimony⁴) that what has gone wrong is that there is a defect in the person using the justification, for example, a failure to make their own decisions and judgements. But consider the following case. Suppose we are having a discussion and someone we know to be wise, and reliable on this issue, claims that it is good when people spend their lives engaging in poverty relief. We ask him what the justification is for that. He replies that he has just given us a justification, namely his testimony.

This is an unreasonable or infelicitous response. When we were looking for a justification, we were looking for something further, something which isn't

⁴ E.g. Nickel (2001), Hills (2009), Crisp (2014)

provided merely by his providing testimony and giving us reliable (though defeasible) evidence that it really is good when people spend their lives engaging in poverty relief. Here the inadequacy of the justification provided by testimony cannot be traced back to a defect in a person using the justification, for, not having yet taken up the testimony, we are not yet using the justification, so there is no defective person. Therefore the testimony must be intrinsically unfit to be a justification, even if it reliably indicates the truth. There would not merely be something amiss with us if we took up this testimony; the testimony itself is amiss as a form of justification.

A clarification of the above argument. When a wise person tells us that it's good when people spend their lives engaging in poverty relief, this doesn't work as a justification of that claim. But it does work as a (perhaps partial) justification of the claim that we ought to believe that it's good when people spend their lives engaging in poverty relief. It's good evidence that that claim is true; supposing that it's important enough to bother remembering, which it is, then it is indeed true that we ought to believe that it's good when people spend their lives engaging in poverty relief, for the reason of the wise person's testimony. But it is important not to confuse justification of belief with a justification of what is believed.

These cases mark a contrast between the normative and the non-normative. Take the non-normative analogue of the first kind of case. If someone believes that *Hovertank 3D* is the first game to use ray-casting technology, and their justification for this claim is that they were told so by someone who reliably knew about the issue, there is no problem with this, no aspect in which it is inadequate as a justification. As for the second kind of case: suppose we are wondering whether *Hovertank 3D* is really the first game to incorporate ray-casting technology. Someone we know to be reliable on this issue tells us that it is. This is a perfectly adequate justification of the claim that *Hovertank 3D* is the first game to incorporate ray-casting technology. If one were to ask the giver of testimony for a justification or proof of that, it would not be obviously unreasonable for them to reply that they had just been given a justification: the testifier says so, and *ex hypothesi* they are reliable

on this issue. So truth indication appears to suffice in such cases as the relevant norm of a good justification.

We may still be inclined to think that whatever peculiarity is going on here, it is one confined to the case of moral testimony, and therefore must be traced back to some feature of testimony itself. But we can broaden beyond the case of testimony, to a wider class of evidential connection that I shall call ‘incidental connections’.

Suppose we modify the above scenario to one of counter-testimony. Someone we know to be reliably wrong tells us that it’s good when people spend their lives engaging in poverty relief. We can’t make sense of why should be wrong so consistently. Still, we find that that’s what they are, and we conclude on inductive grounds that they are a reliable source of indirect information on the truth. We might infer from their remark that it’s probably not good at all when people spend their lives in pursuit of it.

Again, though, this is not a justification of the claim that it’s good when people spend their lives engaging in poverty relief. If someone presented with this reliable counter-testimony then asked what was the justification for the claim that it’s not good when people spend their lives engaging in poverty relief, it would be unreasonable or infelicitous for someone to retort that they had just been given a justification. It’s not just counter-testimony, either: any incidental connection would be inadequate as a justification.

Suppose I know from experience that good actions tend to be those Ted happens to enjoy doing. Again I’m not sure why Ted should enjoy doing them; I know for sure that it’s not because Ted thinks that they are the right thing to do and enjoys doing the right thing. Rather, Ted just seems to enjoy them peculiarly. So I conclude on inductive grounds that when I see Ted enjoying doing something that I should conclude, or at least suspect, that it’s a good action. Then I see Ted spend his whole adult life engaging in poverty relief with that peculiar enjoyment he gets when he does something right. I conclude that it is good when people spend their lives engaging in poverty relief. But I don’t yet have a justification for that claim merely by having that evidence.

Again, both of these contrast with the non-normative case. If someone is reliably wrong as to who directed what film, then they may be an indirect source of information as to who did direct which films. Apart from the unreliability of this evidence, there is no further problem. And were Ted to enjoy doing things which, it turned out, were likely to generate a lot of income for Ted, then again I could have a justification of what was financially best for Ted based on what Ted actually does; and apart from the unreliability of this inductive evidence, there is no intrinsic deficiency with this as evidence.

What was missing in the above cases that accounts for the extra element in justification of normative propositions, besides truth indication? Return to the case when a wise person, when asked for a justification of the claim that it is good when people spend their lives engaging in poverty relief, replies that they have just given us a justification, namely their testimony. This was inadequate. What this shows is that truth indication, as the criterion of a justification for normative propositions, is misguided. Rather, when we want a justification, we want the reasons why it's true, if it is, that that it's good when people spend their lives engaging in poverty relief. Call this the Identity Thesis.

Moral testimony can be an appropriate part of a justification for a normative claim in some circumstances, particularly those where there is some constraint on reason-gathering. Say one is in a case where one has only a very limited amount of time to decide what one should do. In these circumstances, it is legitimate, even right, to take up someone's testimony as to what one should do, assuming they are reliable. In such a case, it can be a justification for what someone does that they were told reliably that they ought to have done it. This does not refute my view, but confirms it: in such a case, it is *because* they were told reliably that they ought to have done it that they had sufficient, even conclusive, reason to do it. Again, the justification displays the reasons why there was sufficient or conclusive reason for them to do it. So this point should not undermine the earlier thought experiment or the Identity Thesis.

What we want when we ask for a justification of the claim that it's good when people spend their lives engaging in poverty relief is not only for an indication of the truth of that claim. What we want is an account of the grounds of that truth, if it's true—grounds without which it would not be true, grounds which explain why it is true. Only this makes for a true justification of a normative proposition. In non-normative cases it obviously sometimes happens that we want an explanation of a truth: why did the Colonel commit the murder? But to ask for an explanation of this is not to ask for a justification for its truth, and vice versa. What marks out normative propositions as distinctive is that, for them, a justification is an explanation.

An example to illustrate the Identity Thesis. Why is perfidy absolutely forbidden? Because any possibility that a surrender has perfidious intent makes it rational for the enemy to have a take-no-prisoners policy. It undermines a system of surrender worth keeping. This reason justifies the claim that perfidy should be absolutely forbidden. The reason indicates that the claim is true. But it also gives the grounds for that truth in giving reasons for it. The normative claim is conceived as true for those reasons. Perfidy is forbidden for the reason that it undermines surrender, because it undermines surrender.

My argument above that justifications are best conceived of as explanations relied on two intuition pumps. Intuition pumps are a very risky kind of argument. Apart from the risk that the reader will simply not share those intuitions, there is a danger that even if they share those intuitions, perhaps they do so for the wrong reasons. Perhaps when it is unreasonable for someone to characterize their own moral testimony as a justification for some normative claim to someone looking for a justification of that claim, this is really just a condition of conversational relevance. Someone who receives moral testimony and then asks for additional justification is wanting justification that goes beyond that testimony, so the testimony wouldn't qualify as a relevant kind of justification. This need not imply that it's constitutive of justification that testimony or counter-testimony is excluded.

I therefore offer an additional argument here as to why both testimony and more general incidental connections should fundamentally fail as justifications. Regarding testimony, someone who gives testimony that one ought to ϕ requires a justification before themselves that one genuinely ought to ϕ (if that claim requires justification at all), and clearly their own testimony cannot be part of that justification. So there is a sense in which, even if a wise person testifies that one ought to ϕ , that just staves off the question of what the ultimate justification for that is. The justification provided by testimony is parasitic on the ultimate justification, and in that sense it is deficient as a justification: it doesn't give us the real reasons why one ought to ϕ , or why some x is good, etc.

A structurally similar argument can be offered regarding the evidence provided by incidental connections, such as a reliable connection between what a particular person says about norms and the truth of the opposite of what they say. If it is a genuinely incidental connection, then if there is at all a justification for the truth of the opposite of what they say, it is a justification independent of them saying it. So again there is a sense in which the fact that there is a reliable connection between what they say and the truth of the opposite just staves off the question of what the justification for that truth is. The ultimate justification is knowable independently. That the ignorant person says the opposite doesn't give us the real reasons; it is deficient as a justification.

Why don't both of these arguments carry over to the case of non-normative justification, with false consequences? After all, it's not as though someone's testimony that the Colonel is the culprit would be a poor or deficient justification. But here they might have a justification which they cannot share with the person they give testimony to. Perhaps they saw the Colonel commit the murder. The closest other people can come to this justification is through their testimony that they saw. The justification the testifier has may be unshareable. So the justification for the Colonel's guilt which consists in their testimony may be the best justification there is; it doesn't stave off the question of what the ultimate justification is. With justifications for normative propositions, in contrast, it is difficult to conceive of a justification which is

unshareable in this way. I can't think of any examples, anyway: if the reason it is wrong to prank the insecure is because of the torment it causes, that is a justification one can share directly.

Regarding the case of an incidental connection in the case of a non-normative justification, things are similar. Suppose that a particular detective is consistently wrong about who the culprit is in the cases they investigate. It may turn out that the killer is so competent that they have got rid of every other scrap of evidence. There may be no more evidence for their being the killer other than the sheer fact that a particular detective known to be reliably wrong has denied it. (Admittedly this wouldn't pass the test of reasonable doubt). So it may not be the case that a better justification is being staved off. The incidental connection may be the best evidence we have. So it cannot be fundamentally deficient on the grounds that there is always a better, more ultimate justification. For propositions with a normative subject-matter, in contrast, there is no analogue of disappearing evidence. If the reason it is wrong to prank the insecure is because of the torment it causes, there is no sense to made of the possibility that that reason might just disappear and no longer be useable as evidence (leaving aside the separate possibility, of course, that the reason might turn out to be a bad reason, or based on a false assumption).

The Identity Thesis says that for a normative proposition to be (ultimately) justified just is for its truth to be explained. To have a (ultimate) justification for a normative proposition just is to have a putative explanation of its truth, to be able to give grounds for its truth. More, this is peculiar to normative propositions; the identity of justification and explanation of truth does not hold for the non-normative case or the case of action. This claim needs clarifying in some more ways.

First: justifications give the grounds of a normative truth. But what is the notion of grounding here? It is that required by explanation. What I claim fundamentally amounts to this: for normative propositions, an adequate justification of them is an explanation of why they hold. So to give the grounds of a normative truth is to explain why it holds.

We can conceive grounds on an analogy with causes⁵, and moreover, as species of a common genus. If *c* is the cause of *e*, then *c* is something because of which *e* happens. *E* depends on *c*. We can think of causes proper (and grounds as well) as one species of a more general genus of dependency relations. Many dependency relations are not causal. If Xanthippe's being a widow at t_2 depends on Socrates' having died at t_1 , it is notorious that this is not causation proper⁶. There are other constitutive dependency relations: a conductor's being hot depends on its particles having a high average level of kinetic energy. This is not causal either. We should see normative dependency relations as just one more type of non-causal dependency relation.

Second, the equation of justification and explanation does not entail that one is only justified in believing a normative proposition just in case one has a justification for it. That this is not so is shown by the case of moral testimony used above. There, one is justified in believing the relevant normative proposition, that it is good when people spend their lives engaging in poverty relief on the basis of a wise person's testimony to that effect. But this is not to have a justification of that proposition.

Another reason justification of belief in a normative claim comes apart from justification of a normative claim stems from the possibility that some normative truths are intuitively knowable, delivered by a moral sense or something similar. If there is a moral sense, and I intend to be neutral on this, then one may be justified in believing the deliverances of that moral sense; it's reliable and so on. But even if there were a moral sense, then one would not have a justification of the deliverances of that sense merely by being subject to them. To have a justification of those deliverances, one would in addition need to formulate grounds for why the deliverances of one's moral sense should in fact hold.

Thirdly, this claim needs to be distinguished from a different claim that is sometimes made in the philosophy of action, for example by Davidson⁷, that

⁵ I owe much of this section to the discussion in chapter 7 of Ruben (1990).

⁶ It is often called 'Cambridge determination': see, for example, Ruben (1988).

⁷ See Davidson (1963)

the justification (or ‘rationalization’, the word which Davidson uses) of an action is what explains it, because the reasons for which an agent performs an action are the causes of his action. If this is true, then this is a sense in which justification is explanation. But this is not what I mean when I claim that justification is explanation, because I am talking about the justification and explanation of normative propositions and truths rather than actions. To justify a normative proposition is to explain why it is true, but to justify an action is not to explain why it is true; that’s meaningless. Hence the importance of distinguishing.

Fourthly, the status of the identity of justification and explanation needs clarifying. It is not prescriptive in the sense that when people justify normative claims, what they *should* be doing is trying to give an explanation of the truth of those claims. Nor is it descriptive in the sense that when anyone gives a justification for a normative proposition, what they are in fact doing is trying to give an explanation of its truth. Rather, the claim is that if anyone has a justification for a normative proposition, what they have is an explanation of its truth. What it is for a normative proposition to be adequately justified is for its truth to be explained.

Earlier I claimed that justification requires reasons. But what is a reason? After making some general points about reasons, I will set up an explanatory challenge. Resolving this challenge, I argue, is a requirement on any adequate account of reasons.

Every reason counts in favour of a particular claim’s being true. That a reason ‘counts in favour’ of a claim is one of many locutions that have roughly the same meaning. We say that *X* is a reason ‘for’ a conclusion or that it ‘weighs in favour’ of that conclusion; that in a list of reasons some belong in the ‘pro’ category with respect to a particular conclusion, as opposed to belonging in a ‘con’ category; and some have said that a reason to ϕ entails that one ‘pro tanto ought’ to ϕ , meaning that as far as that reason goes, one ought to ϕ .

Reasons of this kind figure in justifications. To justify the claim that it’s wrong to prank the insecure, it’s necessary to adduce reasons. One might bring to bear the likelihood of upsetting them, the likelihood that they will

fail to recover quickly, and the likelihood that given the above any prank will fail to be funny. The above three considerations serve as reasons which collectively justify the claim that it's wrong to prank the insecure.

Any reason is based on a fact which 'serves as' a reason. In the above example, it's the fact that pranks will likely upset the insecure that serves as a reason for not pranking them. For the reason to be valid, the fact on which it is based must indeed be a fact. No false claim can serve as a valid reason. If it were simply false that the insecure are unlikely to find pranks funny, then that would not be a genuine reason for it's being wrong to prank them.

To say that reasons which count in favour of a particular claim are the basic elements of a justification for that claim does not yet settle the question of how the justification is to proceed. A particular problem arises over cases where there are reasons which count in favour of competing claims. In virtue of what should the reasons, taken as a whole, recommend one and only one as the proper verdict? We cannot explain why simply by saying there were reasons in favour of that verdict, for there were reasons against it too. We had this scenario in the case of pranking the insecure. Although there were reasons both for and against that claim, it was reasonably clear that pranking the insecure is wrong (other things being equal, of course). But why should it be? It seems there must be some other feature, in addition to the reasons, that explains why things should fall out in favour of one verdict but not the other. What is this feature? Any adequate account of the relation between reasons and justification must say what. Furthermore, in light of the Identity Thesis, this feature of justification is thereby a feature of explanation of normative truth. Why, in scenarios of conflicting reasons, things should fall out one way rather than another, then becomes specifically an *explanatory* challenge—one which tries to say *why* a particular normative truth should hold rather than a competing one, on the basis of the reasons.

One popular kind of view for answering this explanatory challenge is specified in terms of a metaphor of weightiness. Reasons weigh in favour of claims, conflicting reasons weigh against each other, and all the reasons taken together collectively weigh in favour of one claim or the other or none at all.

Roughly, whichever of the competing propositions have the weightiest set of reasons in favour of it, that proposition is justified by the reasons (where there is no weightiest set of reasons, no verdict is justified). There are two principal versions of this story which I will now describe.

One is the ‘balance of reasons’⁸. According to this ‘absolutist’ conception, every reason is associated with a certain ‘weight’: a reason’s weight is one and the same as the extent to which it counts in favour of a particular claim’s being true. We can assess the weight of the overall case for a particular claim’s being true by aggregating the weights of the reasons which count in favour of that claim’s being true. After doing the same for all the competing alternatives, we are to select the weightiest overall case. Roughly, the conclusion of the weightiest overall case is the justified conclusion⁹. In the example above, the reasons in favour of not pranking the insecure, that it is likely to upset them and not to be funny, outweigh the reason against, that it represents its own peculiar form of exclusion.

The second conception, which I shall call the ‘relationalist’ conception, denies that every reason is associated with a certain weight, and instead affirms that reasons bear relations of *being weightier than* to other reasons. This relation is supposed to be transitive, irreflexive and asymmetric, and sometimes to be complete¹⁰. On this view, crudely, a claim is justified just in

⁸ The balance of reasons conception is explicitly endorsed by Broome: the metaphor of balance is introduced in his discussion of ‘weighing explanations’ on p.52 of his *Rationality Through Reasoning* (2013), and defended throughout the rest of that chapter. It is a little less explicit in Skorupski. Skorupski lays out an absolutist (in my sense) conception of the strength of reasons on p.36, but doesn’t appear to mention balance. The metaphors it invokes are pervasive outside philosophy. John Horty quotes (see p.3 of Horty (2012)) correspondence from Benjamin Franklin which indicates that he may have not only endorsed this conception, but used it in everyday life (though it may be ambiguous between the balance of reasons conception and the relationalist conception).

⁹ Subtleties should be introduced when dealing with cases where overall there is some reason to do X, but not enough. See p.38 of Skorupski (2010)

¹⁰ For instance, see Broome’s attack on the possibility of incommensurability in his *Weighing Lives* (2004).

case the set of reasons in its favour is weightier than the set of reasons in favour of each and every competing claim.¹¹

The weighing model needs to answer the explanatory challenge which arises concerning in virtue of what one of two competing normative propositions for which we have reasons should count as the proper conclusion to draw from those reasons. Its answer must be: it is in virtue of one set of reasons being the weightiest set. Weight thereby figures as an explanatory notion in this view¹². At first pass there seems something unsatisfactory about this conception: uninformative and empty.

Moreover, the weighing model sheds no light on when justification is needed. Rather, it takes for granted ideas of when justification is needed: when there are competing claims to decide between which we need to go to the reasons. *Every* normative proposition competes in this sense with its own negation in that one cannot accept both. But I do not think it is obvious that every competition in this sense needs justification, that every normative proposition

¹¹ This second conception is explicitly endorsed by Schroeder: “I will define the weight of a set of reasons not by establishing a correspondence between sets of reasons and some further thing—amounts of weight—but by defining a partial ordering on sets of reasons: the *weightier-than* ordering.” (p.126-7 of *Slaves of the Passions* (2007)). John Harty’s application (see especially chapters 1 and 2 of *Reasons as Defaults* (2012)) of default logic to the structure of justifications is arguably a notational variant of this conception as I argue in the appendix to a later chapter. Finally, something like this conception has been recently endorsed by Scanlon¹¹ though he rejects the correlative claim that *being weightier than* is a logically primitive relation between reasons. According to Scanlon, reasons bear *weightier than* relations to other sets of reasons, but all this literally means is that the weightier set is one from which we can draw the justified conclusion; here, there is no claim that *weightier than* relations play any role in determining the justification. (See the definition of ‘outweighing’ on p.108 of *Being Realistic about Reasons* (2014). Also on p.111 he says ‘The strength of a reason is an essentially comparative notion, understood only in relation to other particular reasons.’)

¹² However, one might also think with Scanlon that weightiness is not an explanatory notion so much as an accountancy one. To have a certain weight property simply *means* that if it is the weightiest the conclusion it endorses is correct. Because it’s defining of weights, on this view, there is no sense in which weights explain why a certain conclusion is correct. To do that, what it is to be of a certain weight has to be defined separately from the consequences of having weights. Scanlon’s account of weights therefore is not apt to answer the explanatory challenge.

(if it is true) needs a justification as to why it should be true over its contradictory.

This leaves us with a task. To what extent can the weighing model answer the explanatory challenge? How does it fail, if at all? I will explore these questions in the next chapter. I also explore in more detail the assumptions of the explanatory challenge. I argue that the origin of the explanatory challenge lies in the work of David Ross, and that Ross' particular commitments make the weighing model seem natural. The work in the chapters after that is partially a refashioning of these commitments. In the next chapter I also criticize the weighing model's answer to the explanatory challenge, and argue that we should look elsewhere.

The Weighing Model: Origins, Motivations and Problems

According to the weighing model, what determines the justified conclusion when there are competing reasons in play are the weight properties of those reasons and their relations to each other. The justification for the right conclusion consists in the reasons and their weight properties concatenated correctly. This characterization is incredibly abstract, so in this chapter I will make it less abstract by discussing each version of this view separately. I do not have a master argument against this family of views, though I do have a debunking explanation of why this family of views exists. So, first, I'll motivate a certain way of thinking about the issue of justification within which the above family of views seems attractive, and draw a moral about why they are inevitably flawed. Then I'll discuss each one in turn: what should emerge is a series of problematic, unsatisfactory views. So this should suffice to motivate the rejection of the entire family.

W.D. Ross arguably put forward the first view of this kind in his 1930 book *The Right and The Good*. Several important points emerge from a consideration of this view, which is why it will occupy us for a while. The discussion is important in a variety of ways. Firstly, Ross lucidly set up a genuine problem which it is helpful to regard the weighing model as trying to solve: so it gives us a standard by which to evaluate the success of those views. Whenever those views appear in print, they are not thoroughly motivated, so it is to Ross that we must turn to find a suitable standard of evaluation. It is also a problem I expect my view to be able to solve, so it is important on that account also. This problem is precisely a version of the explanatory challenge I set up last chapter. The discussion of Ross serves two purposes, then. First, it clarifies the explanatory challenge and motivates it more thoroughly. It turns out that there are two main ideas underlying the challenge: the idea that reasons form an irreducible plurality which cannot be cashed out in terms of a common kind, and the idea that conflicts between reasons are rationally, or at least determinately, resolvable. Secondly, the way Ross seeks to establish these claims itself leads on naturally, if not inevitably,

to the weighing model, even though Ross himself (as I will show) was somewhat ambivalent in this respect. Moreover, the sheer stature of Ross makes it probable that his discussion was widely read and his ideas influential: so Ross is potentially of genealogical interest, especially with reference to how close he comes to endorsing some kind of weights framework.

Ross' concern with the structure of justification arises from his sense of inadequacy with the utilitarianism of Moore and Rashdall,¹³ according to which 'all 'conflicts of duties' should be resolved by asking 'By what action will most good be produced?'.¹⁴ Ross thinks that if an act is to be one's duty then its dutifulness must be grounded in something other than in its tendency to promote the most good. Even if it is one's duty to perform a certain act because it produces the most good, that is only because the action is right *qua* its property of having the tendency to produce the most good, not because of the results of the act:

'That which is right is right not because it is an act, one thing, which will produce another thing, an increase of the general welfare, but because it is itself the producing of an increase in the general welfare...we have to recognize the *intrinsic* rightness of a certain type of act, not depending on its consequences but on its own nature.'¹⁵

So Ross considers the various descriptions under which a particular act may fall. An act under some of those descriptions may be wrong, under others right (other things being equal): an act under the description 'breaking a promise' is intrinsically wrong, under the description 'promoting the most benefit to others' intrinsically right. These because, respectively, we have a general duty to keep promises, and a general duty to promote the most good. This thought puts Ross in a position to articulate the real source of his dissatisfaction with utilitarianism:

'[Utilitarianism] says, in effect, that the only morally significant relation in which my neighbours stand to me is that of being possible

¹³ P.18 of *The Right and the Good* (1930)

¹⁴ P.19, *ibid.*

¹⁵ P.47, *ibid.* Ross is avowedly influenced by Prichard's 'Does Moral Philosophy Rest on a Mistake?' (1912), which receives an explicit mention in the preface. The thought that oughts can be derived only from oughts and not from the good is articulated in that paper.

beneficiaries of my action....They may also stand to me in the relation of promisee to promiser, of creditor to debtor, of wife to husband....and the like; and each of these relations is the foundation of a *prima facie* duty, which is more or less incumbent on me according to the circumstances of the case.’¹⁶

Ross clearly thinks that the fundamental ethical thoughts utilitarianism admits are too few: acts are nonderivatively right under more descriptions than merely ‘promoting the most benefit to all’. For instance, one might reasonably think of an act that it is right since it means keeping a promise, without having reached any conclusions about what good it will bring. This is key to Ross’ subsequent criticism of utilitarianism: for Ross, utilitarianism is refuted precisely by the fact that sometimes what makes an act *prima facie* right does not make it *prima facie* likely to bring about more good than the alternatives¹⁷, even if it is actually right. What a consideration of the *prima facie* makes clear is that rightness and optimality are not equivalent, as Ross’ utilitarian claims.

Ross claims that our actual duty is derivative on our *prima facie* duties. ‘Whether an act is a duty proper or actual duty depends on *all* the morally significant kinds it is an instance of.’¹⁸ Under some of these *prima facie* duties, some performable actions will be right, under others wrong. This gives us a sense in which *prima facie* duties conflict: they favour competing verdicts. But actual duties exist nonetheless. So the question is left open as to how to resolve conflicts between *prima facie* duties.

A general argument can perhaps be drawn from Ross’ remarks on utilitarianism. A principle for resolving conflicts between *prima facie* duties would have to take the form ‘It is right that we ϕ ’, where ϕ stands for an action under some general description e.g. promote the most benefit to all. But any such principle is simply equivalent to saying that it is always our duty to ϕ , which means that it cannot coherently appear that it is ever our duty to do something other than ϕ . It then follows that there can be no *prima facie* duty to do something other than ϕ , since a *prima facie* duty is the kind of duty that is apt to be our actual duty. So any principle for resolving conflicts

¹⁶ P.19, *ibid.*

¹⁷ P.36, *ibid.*

¹⁸ P.20, *ibid.*

between *prima facie* duties in favour of one particular duty effectively rules out the other *prima facie* duties from being *prima facie* duties at all. And this is inconsistent with the facts of moral experience. This argument is important for motivating the explanatory challenge. Its conclusion is the irreducible plurality of reasons, which is a central plank in the explanatory challenge.

So, for Ross, one resolves conflicts between *prima facie* duties not with reference to a third thing like a project or principle. The key move is the denial that *prima facie* duties are of a kind, or measurable in terms of a common kind like how much good each alternative produces. One resolves conflicts between *prima facie* duties by looking at those duties themselves. And this leads to an important point about the way Ross thinks of *prima facie* duties. A *prima facie* duty is not just something which initially appears to be one's duty but which might turn out not to be one's actual duty, as for instance when it initially appears to us to be our duty to keep a promise but then, on reflection, we ought actually to break it. It is that, but it is more. A *prima facie* duty, for Ross, is an actual duty unless there is some other *prima facie* duty with which it comes into conflict. As Ross puts it:

‘I suggest ‘*prima facie* duty’...as a brief way of referring to the characteristic...which an act has, in virtue of being of a certain kind..., of being an act which would be a duty proper if it were not at the same time of another kind which is morally significant.’¹⁹

Ross is remarkably ambivalent on whether we arrive at the correct conclusion by rationally putting together the weights of the *prima facie* duties. A couple of remarks suggest that he would favour the view that when *prima facie* duties conflict, one of those becomes our actual duty by being a weightier duty:

‘It may be said that besides the duty of fulfilling promises I have and recognize a duty of relieving distress, and when I think it right to do the latter at the cost of not doing the former, it is not because I think I shall produce more good thereby but because I think it is the duty which is in the circumstances more of a duty.’²⁰

And elsewhere:

‘Every act therefore, viewed in some aspects, will be *prima facie* right, and viewed in others, *prima facie* wrong, and right acts can be

¹⁹ P.19, *ibid.*

²⁰ P.18, *ibid.*

distinguished from wrong acts only as being those which, of all those possible for the agent in the circumstances, have the greatest balance of *prima facie* rightness, in those respects in which they are *prima facie* right, over their *prima facie* wrongness, in those respects in which they are *prima facie* wrong...for the estimation of the comparative stringency of these *prima facie* obligations no general rules can, as far as I can see, be laid down.’²¹

The language of ‘more of a duty’, the ‘comparative stringency’ of duties, together with the classic metaphor of ‘balance’, express his attraction towards a weighing model. The other side of the ambivalence is his denial that such properties are epistemologically accessible in the same way as is the existence of the *prima facie* duties. This is plain in the last sentence of the above quote. His positive remarks on the epistemology of actual duty are vague: he suggests, quoting Aristotle *Nic. Eth.* 1109 b 23, 1126 b 4, that ‘the decision rests with perception.’ He also refers to a ‘sense’ of our particular duty, ‘preceded and informed by the fullest reflection.’²² Finally, he suggests that when in a situation of conflicting duties, ‘what I have to do is to study the situation as fully as I can until I form the considered opinion (it is never more) than in the circumstances one of [the *prima facie* duties] is more incumbent than any other.’²³ Perception, reflection or studying the world: what kinds of ways of gaining knowledge do these rule out (apart from, perhaps, testimony)?²⁴

The picture that emerges from this exposition is of a view which is motivated not by positive argument or a thorough study of examples but primarily by what it wants to reject: views which attempt ‘to state a single characteristic of right action’²⁵, and utilitarianism as a dominant view of this kind. Moved by the thought that saving the *prima facie* appearances requires rejecting the

²¹ P.41, *ibid.*

²² P.42, *ibid.*

²³ P.19, *ibid.*

²⁴ His negative remarks on the epistemology of our actual duty are more determinate. He denies that our actual duty is self-evident when there is a conflict of *prima facie* duties: in conflict scenarios, ‘we are (I think) well aware that we are not certain whether we ought or ought not to do it; that whether we do it or not, we are taking a moral risk.’ (P.30, *ibid.*) And he denies that knowledge of our actual duty could ever be deductive: ‘the only possible premises would be the general principles stating [acts’] *prima facie* rightness or wrongness...there is no principle by which we can draw the conclusion that it is on the whole right or on the whole wrong.’ (P.31, *ibid.*)

²⁵ P.16, *ibid.*

general principles (in the light of which such appearances must inevitably be seen as illusory), Ross chooses to reject the general principles in order to save the appearances. The result is that when trying to account for the ground of actual, all-things-considered duty, there is not merely a missing bit of theory but active confusion over that in virtue of which actual duties hold. Some duties are more than others, but not so obviously so as to prohibit the possibility of the lesser duties' appearing to be duties at all. So the facts of right action must be non-obvious: informed by thorough reflection and study of a kind that apparently cannot be spelled out on pain of being obvious in this objectionable way.

This particular result follows from Ross' concern to protect the genuineness of the *prima facie* duties. Translating Ross' general view into a modern idiom makes clear the way in which this thought is mistaken. The *prima facie* duties are translated into 'reasons'. Instead of an actions' rightness being determined by the sum of the ways in which it is *prima facie* right or wrong, it is determined by all of the reasons for its being right or wrong taken together. For Ross, a duty could be our actual duty by being 'more of a duty' than all of the other pertinent duties; in the modern idiom, a reason for an act's being right establishes its rightness just in case it is weightier than all the reasons against, or the set of reasons for that act's being right is weightier. But what is clearly expressible in the modern idiom is the difference between a reason and what gives rise to that reason. A concern to keep one's promises or the value of keeping them may give rise in any situation to a reason to keep that promise. A reason to promote the most benefit to others may clash with and be a better reason than the reason to keep one's promises, but it does not follow from this that the concern to keep one's promises is not appropriate or that we should not be so concerned. It did follow for Ross, in contrast, that if a duty to do X is clearly stronger than a duty to do Y, then with this knowledge the duty to do Y cannot coherently appear to be a duty in the first place.

Because the modern idiom does not produce a similar result, there is no barrier within it to taking conflicts between reasons to be resolvable in a self-evident or deductive way, which is what Ross denied. If reasons have weights, then one can simply add them together and derive the result. This kind of view

is more rationalist than Ross, in the relatively minimal sense that there is something we realize when we realize why an act is overall right, and the act's rightness deductively follows.

Despite the differences resulting from this reformulation, an examination of Ross is valuable not merely from a genealogical perspective but also because Ross did latch onto a genuine problem, a problem which the introduction of weights appears to solve. This problem, as I have argued, arises not just for Ross but for all parties, and so it is incumbent upon us to answer it. One of Ross' concerns was to preserve a sense in which the conflicts between reasons is determinately resolvable. The other is to preserve the sense in which reasons, or *prima facie* duties, genuinely conflict. This conflict cannot be preserved under a theory which espouses a single characteristic of right action, such as utilitarianism. We can get this result even without thinking of reasons as *prima facie* establishing what we must do.

Suppose we are in a choice situation and we are determined to do what promotes the most good to others. There are two choices, A and B: A would bring about x good, B would bring about y good, and $x > y$. Could we construe this as a scenario where the possibility of bringing about x good is a reason to choose A, the possibility of bringing about y good is a reason to choose B, and the general dictum to promote the most good adjudicates between these? Perhaps, but this characterization is not forced upon us. We could equally construe the choice as requiring a simple syllogism. The major premise: whichever choice would bring about the most good should be chosen. The minor premise: choice A brings about the most good. Conclusion: A should be chosen. Characterized in this way, there is no conflict between reasons. So this is the respect in which the endorsement of a single characteristic of right action fails to preserve a sense of conflict between reasons: it does not force upon us a characterization of justificatory questions as a conflict between reasons, since a characterization in terms of a syllogism serves equally well.

The second ingredient in Ross' problem was that conflicts between reasons are resolvable (at least for the most part) in that whenever there is an ethical question with reasons both for and against a central proposition, there is a

correct answer as to whether that proposition is true or false, and its truth or falsity is satisfactorily explained by the reasons. For example, if we are considering whether home schooling is a good thing, there are reasons both for and against this. In some circumstances it can raise standards, it allows more flexibility and the possibility of an education tailored to the needs of the individual, but also it can inhibit the development of a child's social skills. Even here, where the issue is somewhat difficult to decide, there are one can lean one way or the other. I myself am inclined to the view that home schooling is not a good thing, though I acknowledge that there are genuine reasons in favour of it.

With these two points in hand, we can formulate Ross' problem as an explanatory challenge. Take the proposition: home schooling is a good thing. Call this P, and call the proposition that home schooling is a bad thing Q. As we've seen, there are reasons for P, and reasons against P and for Q. P and Q aren't exhaustive: it could be that home schooling is ethically neutral. Suppose, with me, that Q is true: home schooling is a bad thing. We might think that this is because of the reasons in favour of Q. But there are reasons in favour of P as well. So this doesn't explain why the fact that there are reasons in favour of Q should make Q true: there are reasons in favour of P, and yet those don't make P true. So there is an explanatory challenge: why should the reasons in favour of Q establish Q, and yet the reasons in favour of P fail to establish P?

To answer this challenge, some additional structural feature is needed. Weights are a good first pass at this. If the reasons in favour of Q are weightier than those in favour of P, then it makes sense that they should establish Q while the reason in favour of P fail to establish P. We could then extend this into an argument for the existence of weights, on the grounds of an inference to the best explanation. It is not yet an inference to the best explanation because we haven't surveyed other good explanations. In this thesis I aim to introduce one; part of its purpose is to disrupt this implicit argument to weights from the consideration of the best explanation of the puzzling phenomenon noted above.

The first point of interest with the weighing model is not an objection so much as a challenge designed to clarify the commitments of the view. From this challenge it becomes clear that weight-based views gain apparent traction solely on an abstract front. We should evaluate them solely by their ability to solve Ross' problem, if the argument below is correct.

A crucial aspect of this view which was also arguably present for Ross is that whether a reason has a certain weight, or whether a duty is more of a duty than other duties, is itself a normative question. In the same way that someone is making a normative claim when they say that the fact of making a promise is a reason to keep it, they are making a normative claim when they say that this reason is weaker or stronger than the reason to promote the most good, in scenarios where they come into conflict. Schroeder is explicit about this²⁶, but Broome is not²⁷. In any case, whether or not writers acknowledge it, it is a straightforward consequence of the view. That a reason is weightier than other reasons involves some evaluation or appraisal, namely that that reason is one that we should take more seriously than other reasons, or should 'place more weight on it'²⁸.

As to the first challenge: it also is rooted in the awkwardness of the dissociation between the weight of a reason and the reason's existence. I suggest that weights gain traction as an attempt to solve Ross' problem outlined above. But aside from this, weights are superfluous. They are not superfluous in the sense that they are an idle wheel. Rather, they are superfluous in the sense that they do not shed any light on justifications, even though they are not an idle wheel. There are a variety of ways in which a theory may shed light on justifications, each of which I will consider in turn: it may shed light on what justifications are, how we know them, when they are appropriate, or how we think about them. I will argue that weights do none of these things.

²⁶ See the formulation of the 'Attractive Idea' on p.130 of *Slaves of the Passions*

²⁷ In his writings on it in *Rationality Through Reasoning* chapter 4.

²⁸ Schroeder, p.130 of *ibid.*

If reasons have weights, then whenever it is clear what weights the relevant reasons have we can deduce what conclusion is justified. We can contrast this with a weight-less framework in which there is nothing but reasons and the conclusion. It is certainly not clear, whenever we know what the reasons are, we automatically know what the conclusion is. We might know that the potential educational benefit of home schooling is a reason in favour of the worth of doing it, and the potential social inhibition a reason against, without it being clear from that alone whether it is worth doing or not. This was one of Ross's insights: in such a situation we have to consider it more until we come to a considered opinion. So one question to ask of the weights framework is whether it sheds any light on the process of coming to a considered opinion.

The only way it can do that is if the following situation is possible. We know what the reasons are, but we are unsure what opinion we should adopt. Then we consider what weights the reasons have; if we can decide that, we can deduce what the proper conclusion is. What weights the reasons have is clear. So we deduce the proper conclusion. Hence, by adverting to weights, we have resolved our original uncertainty. Weights have done work. The key idea is that the question of what weights the reasons have must be in principle easier to settle than the outright question of what the proper conclusion is. Only if that is sometimes true will weights not be epistemically superfluous.

A consideration of some examples suggests (but does not entail) that this is true. Whenever it is clear what the weights are, it is already clear what the conclusion is. So there is no scenario in which the question of what weights the reasons have is easier to settle than the simple question of what to conclude.

Let's start with the example I've been using. We're wondering whether home schooling is a good thing: against we have the inhibition of social development, in favour of it we have the possibility of tailoring an education to individual needs. Suppose (modifying the example) it turned out that the inhibition to social development was very grave indeed, while the educational benefits were slight. By this point it is fairly clear that home schooling would

be a bad thing (were this true). But it does seem in that case to be more clear that the inhibition of social development is a weightier reason against home schooling than the educational benefits are in favour of it. Weights, then, would appear to add nothing in this example to the consideration of whether home schooling is good or bad.

Here's another example. Someone is burning in a malfunctioning factory and we are the only person around. It is open to us to charge in and try to save them. There is a 50/50 chance of saving them, with the alternative being that we fail and die as well. It's not obvious what we should do (if you still think that we should try to save them, then lower the chance of success until that judgement disappears). But it is not obvious whether the prospect of saving them is a better reason to try to save them than the risk of failure and death is as a reason to refrain from trying. Neither the conclusion nor the weights are clear. Again, weights add nothing to the consideration of the issue. Weights are epistemically superfluous.

Perhaps weights gain traction as a characterization of what it takes for a verdict to be a justified one. If home schooling is bad, then in a parsimonious account, it is bad because it inhibits social development, even though there are educational benefits. In the weights account, it is bad because it inhibits social development, and this reason is weightier than the reason in favour of it posed by the educational benefits. The addition of weights does not seem to add anything to the explanation.

Finally, perhaps weights do work in an account of disagreement. I have in mind something like the following. When Chad and Clive disagree over whether home schooling is good or bad, but agree that the inhibition of social development is bad other things being equal, and agree that the possibility of tailoring an education to individual needs is good other things being equal, but still disagree over whether home schooling is good or bad, maybe we can say what it is they disagree about in terms of weights. Perhaps what it is they disagree about is the weightiness of the various reasons. Chad thinks the inhibition of social development is a weightier consideration than the possibility of tailoring an education; Clive thinks the opposite. It is because

they disagree about this that they disagree on whether home schooling is good or bad. Again, however, it is not clear what this adds to the same scenario redescribed without weights. Perhaps what Chad and Clive disagree about is simply the status of home schooling, even though they agree on what the reasons are. What is more perspicuous about saying that they disagree on how weighty the reasons are?

If weights exist at all, then, we only have reason to think so on the basis of their ability to account, in the abstract, for why the reasons on one side of a case establish the conclusion they favour, whereas those on the opposite side fail. They can do this only if they form part of an adequate explanation for this phenomenon. Even if this were the case, however, they may suffer from an additional problem. Weights must themselves be explained: which weight relations hold between reasons, or how weighty individual reasons are. There is a danger that whatever would explain the weights would explain the original phenomenon the weights were supposed to explain. For example, suppose Chad arrives at the following view. He ought to eschew home schooling because the inhibition of social development is a weightier reason not to do it than is the educational benefit as a reason to do it. It is weightier because social competence is a key ingredient in future happiness, whereas educational disadvantage can be overcome. But we could then say that this difference in whether one can compensate for it itself explains why home schooling is to be avoided (on the basis of the inhibition of social development), even though the potential educational benefit is a reason to do it. What explains the weights can do the work the weights do: it's not clear that anything is left out. Putting this in terms of Ross' problem, we could then say that when a set of reasons R favours P while a set of reasons S favours Q which is incompatible with P , then what explains why P rather than Q (even though there are reasons for both) is itself a further reason. I find this argument from crowding out persuasive. In any case, I will now argue that weights cannot adequately explain the phenomenon in any case.

I argued earlier that the claim that a reason has a certain weight must be normative. If these claims are normative, then they should themselves be subject to justification. It should be reasonable to demand justifications for

the claim that some reason is a weighty reason, and reasonable to expect that such demands can be answered.

A problem arises here for the relational conception of weights. According to this conception, facts about the weightiness of reasons are fundamentally all and only relational facts, so we can talk coherently of one reason or set of reasons being weightier than another, but we cannot talk at this fundamental level of a reason being a weighty one *simpliciter*. So, that a reason is weightier than another one requires justification, and thereby requires justification in terms of reasons.

The problem is best put in terms of an example. It is best to use those examples that the proponents of a particular view themselves use, so I would like to use Schroeder's example of justifying the weightiness of certain reasons. Strangely, however, although Schroeder bases his account on a description of a weightier-than relation, the example he uses only has one reason for belief, and further reasons justifying whether that is a weighty reason for belief. There is no reason it is argued to be weightier-than. So the example does seem to be inconsistent with the thrust of Schroeder's account. Specifically, he gives the example²⁹ of seeing Tom emerge from the library, pull a book out from under his shirt and run off, which gives some reason to believe that he has stolen a book. But how weighty this reason is qualified by how likely it is that we have recognized Tom, for we know that his identical twin brother Tim is also in town. In this example, the reason to believe that Tom has stolen a book is judged for its weightiness without reference to any reason it is weightier than. So we're going to have to modify the example in order to bring out the problem for a relational conception of weights.

Suppose instead that Mary has told us that Tom is a very upright man who would never steal. This gives us some reason to believe that Tom has not stolen the book. We have to weigh this reason against that posed by what we saw: Tom, or a Tom lookalike, run off from the library with a book pulled out from under his shirt. We are wondering whether the latter reason is weightier than the former; if so, we should believe that Tom has stolen a book, and vice

²⁹ P.132, *ibid*.

versa. That the identical twin Tim is in town is a reason to place less weight on what we saw as a reason for believing that Tom stole the book.

We might think, then, that the fact that Tim is in town is a reason to think that what we saw is a less weighty reason than what Mary told us. This has to be true, if the fact that Tim is in town is going to do work in justifying what weightier-than relation there is between these two reasons. But this doesn't quite work as the structure of a justification. It falls foul of what it is to be a reason. To be a reason is, as it was for Ross, to be something that establishes the conclusion it favours unless there are contradictory reasons. So, for instance, that I promised is a reason to keep the promise, and *ipso facto* I should keep the promise unless there are other reasons not to. We can call this the presumptiveness condition. Schroeder also endorses this condition³⁰, as does Scanlon³¹, and Horty³², effectively.

But the above reason, that Tim is in town, falls foul of this condition. To meet this condition, it would have to be true that the fact that Tim is in town establishes that what we saw is a less weighty reason to believe that Tom stole the book than the reason against posed by what Mary told us, unless there are other reasons to believe that it is a weightier reason. But this is false in this case. Suppose there are no other reasons that bear on the matter. Just because Tim is in town, it doesn't follow that what we saw is a less weighty reason than what Mary told us. Perhaps, even though Tim is in town, what we saw is a weightier reason to believe that Tom stole the book than the reason posed by what Mary told us. Hence, it does not make sense to think of the fact that Tim is in town as a reason to believe that what we saw is a less weighty reason than the reason posed by what Mary told us. If it were a reason, it would have

³⁰ See the definition of 'ought' on p.130, *ibid*.

³¹ See the definition of outweighing on p.108 of *Being Realistic About Reasons*.

³² *Reasons as Defaults*, (2012), ch. 1 and 2. A 'default' is a presumptive pattern of inference, so it follows that it establishes its conclusion absent other defaults. Horty also effectively endorses the relational conception of weights in making conflicting defaults resolvable by other defaults. Spelling this out, this would be a case where there is a presumptive pattern of inference A such that when confronted by presumptive patterns of inference B and C such that B and C both pertain presently, then the conclusion is that favoured by B and not C.

to establish its conclusion absent other reasons, but we have seen that it doesn't.

We might try to find a way to escape the presumptiveness condition. Perhaps you can have a reason to believe something without having any reasons not to believe it, and still you wouldn't have to believe it. For instance, suppose you are on a bus and you fleetingly hear someone say that the Russians have just launched a nuclear strike against everywhere other than here. You suspect that you misheard, and it's so implausible anyway. Although you don't have any particular reason not to believe that the Russians have just launched a nuclear strike, you still wouldn't have to believe that they did. So perhaps the condition is wrong. But this case isn't quite right. You do have positive reasons to believe that the Russians haven't just launched a nuclear strike: why would they? What would it get them? So this just isn't an example of a case where you have some reason to believe something and no reasons against believing it. I have not encountered a more plausible case.

Another response to my objection is to suppose that there will always be other reasons that bear on the matter of whether Mary's testimony or our observation is weightier; in that case the scenario in which Tim's being in town is the only reason which bears on the matter of which is weightier is excluded. For this to work, it would have to be the case that there are always other reasons which bear on the matter of whether Mary's testimony or our visual experience is weightier. So we would have to be unable to conceive of a scenario in which there were no additional reasons other than that posed by Tim's being in town. This could be true if those other reasons were bound up with the testimony or observation, so those first-order reasons could not be present without the reasons justifying their weight relation also being present. For example, perhaps Mary's general reliability, in contrast to the general unreliability of fleeting observations, would serve as a reason justifying the proposition that Mary's testimony was weightier than our observation. But that it is not a reason of the right kind; we can conceive of a case where that reason was not present. We could simply suppose that Mary was not generally reliable. It is not clear, then, that there are always other reasons which bear on the relative weightiness of Mary's testimony and other observation, no

matter what the specifics of a scenario. And hence there are scenarios when Tim's being in town is the only additional reason which bears on the matter of the weight relation.

My objection purports to show the inadequacy of the relational-weights framework for making sense of the structure of the reasons in Schroeder's example. We might think that all this shows is that there are certain bits of the framework we should drop (for instance, Ross' presumptiveness condition on being a reason). But I think that my objection brings out a general structural problem which should at least constrain what counts as an adequate reworking of the theory. The challenge for the relational weight theorist is to show that the structure of justification which is to be present in the justification for overall, all-things-considered judgments will be replicated at the level of justifying which weight relations obtain between reasons. The theorist is not entitled to assume this.

This objection doesn't work against an absolute conception of weights, which is consistent with the possibility of the scenario I have used in my objection. According to that conception, a reason's weight is fundamentally a non-relational property of that reason. Although the weightier reasons defeat the less weighty reasons in a justification, what weight is is a non-relational property of reasons. Under this conception, what we justify when we justify the claim that a reason has a certain weight is that it has a non-relational property of having some level of weight w . So we can think of weight-alteration scenarios as follows: one reason S has a certain weight w , the opposing reason T has a certain weight v , and further facts may be a reason R to think that S has a certain weight z , where $z > w$. By the condition and the fact that there are no other reasons, we conclude that S does indeed have a weight of z where $z > w$. It is entirely consistent with this conclusion that $v > z$, that T is weightier than S even with the extra information. And this is a satisfactory reading of the scenario I have used in my objection against the relational conception of weights.

But this absolute conception of weights has its own awkwardnesses in describing the ways in which the weight of a reason is itself justified by other

reasons. As I have said, Broome does not explicitly acknowledge that a reason's having a certain weight is itself normative, so he does not discuss what justifies it. I am extrapolating when I discuss an absolute conception of weights.

If a reason R counts in favour of another reason S's being of weight z instead of weight w , where $z > w$, what determines why it is z that it counts in favour of, rather than v or x , where both v and x are greater than w ? Why one particular level of weight rather than another? In the limiting case, why should R not count in favour of S's having an infinitely large weight, rather than in favour of its being of weight z ? There does not seem to be any ready answer to these questions. So we might try to modify the view: reasons have certain weights. Other reasons may affect that weight by increasing it by a certain amount or decreasing it by a certain amount. R may increase S's weight from w to z , another reason T may decrease S's weight from z to x , and so on, until S has some remaining weight v . But again, as with the objection to the relational theorist: we are not entitled to assume that weights must conform to this structure.

Secondly: given that R counts in favour of S's having a certain weight z , what determines the weight R has in counting in favour of S's having z , the weight which will determine what other reasons will successfully defeat the justification for S's having z ? If the weight of a reason is itself determined by reasons, a need for justification arises for the weight of those reasons which justify a reason's having a certain weight. It is not clear what kinds of reasons would be involved in practice. Suppose the fact that social skills are more fundamental to human development is a reason in favour of the social loss of home schooling outweighing the potential educational benefit. What determines the weight of this reason (call it A)? Perhaps the importance of human development. In turn, what determines the weight of this as a reason in favour of A's having a certain weight? There does not seem to be a ready answer to this question. In general, going back along the chain of justification, on this view, seems likely to lead to justificatory questions which lack an answer. There must be an answer on this view, on pain of the weight of a reason's being unjustified or basic. Perhaps there are basic or unjustified

normative truths, but it is inappropriate to commit to them just because the abstract straitjacket into which one fits the reasons requires it.

The absolute conception of weights is full of abstract pit-holes. It generates needs for justification by its abstract structure alone which have no clear solution. The result is an awkward view: one that does not generate a contradiction but which doesn't fit the reasons as we find them. It is, in virtue of the burdens of justification it generates, a substantive view. This seems to cut against the way in which weight views tend to be presented, namely, as highly intuitive and innocent.

The problems with both the relational and absolute conception of weights stem from the core feature of any weight-based view: the weight of a reason is an extra feature, even separate from, the fact that something counts as a reason in the first place. This separation introduces an awkward dissociation between the reasons behind a reason's having a certain weight, and the basis for the reason in the first place. The separation is in the foreground in Ross: self-evident or deductive principles lie behind something's being a *prima facie* duty in the first place, but nothing illuminating and general can be said about the interaction between these duties.

In this section I have expressed scepticism about weights vis-à-vis their claim to shed light on justifications. But this is what we should expect in light of the character of Ross' view. Ross' view is not motivated by positive argument or a thorough study of examples but because it is a natural alternative to a kind of view Ross was hostile to. It is no wonder, then, that the view developed in this way generates no traction when applied to examples or when evaluated with respect to its ability to deliver positive, interesting predictions. This is the debunking explanation of why weighing models exist: they have been inherited from Ross, but not explicitly, and are therefore treated as the most natural thought, exempt from a need for positive argument. Broome introduces it like this³³ as just 'part of our vocabulary about reasons'. Scanlon introduces it like this³⁴ as simply a datum to be accounted for. This is

³³ P.51, *Rationality Through Reasoning* (2013)

³⁴ P.3, *Being Realistic About Reasons* (2014)

misguided. Such views manage to be substantive in the abstract but superfluous in the concrete; and moreover they are attended with problems even in the abstract. They should be rejected.

Burdens of Justification

In Chapter 1 I claimed that a justification for a normative proposition is best conceived as an explanation of the putative truth of that proposition. The reasons that count in favour of a normative proposition's being true give the grounds on which the putative truth of that proposition depends. For a normative proposition to be justified just is for its truth to be explained.

Spelling out what this means in more detail requires a thoroughly specified notion of explanation. Once we have this notion of explanation, we can give a construal of justification in terms of it. So in this chapter I will specify more thoroughly what explanation is, and then use it to give an account of the nature of justification.

I adopt a particular conception of explanation according to which all explanation is contrastive, in a sense I'll spell out soon. This is not a particularly new view; it has been around since the mid-1970's³⁵. But since its inception it has been concerned with issues in the philosophy of science: for example, clarifying the 'miracles' argument for realism (realism is the best explanation of the success of science, see Van Frassen 1980) or in overcoming some more traditional epistemological problems such as the problem of induction (Lipton 1991). The standard exposition of explanation and its contrastive nature is wrapped up in these concerns. So I aim in this first section to give an exposition of this conception of explanation that is free from these concerns and is more obviously applicable to the domain of the normative. I will play up those aspects of the conception that will be of importance later.

A starting platitude in getting at explanation is that explanation aims at enabling understanding. If we want an explanation as to why *p*, a successful explanation enables us to understand why *p*. This platitude is good for a

³⁵ Van Frassen cites unpublished work by Bengt Hansson (1974): 'Explanations-of-What?' mimeographed and circulated, Stanford University.

couple of observations. First, it does not give us any reason to expect that explanatoriness is usefully formulated in terms of conditions on an *explanans* and conditions on the relation between an *explanans* and the *explanandum*, though it is consistent with this possibility. If explanation's fundamental nature is to enable understanding, we should leave it open at this stage that a good explanation will have lots of separate parts, each tailored to a specific failure of understanding. (As it turns out, this is an aspect of a contrastive theory of explanation). Second, if explanation is there to enable understanding, then we should expect that what makes a good explanation in any particular circumstance depends heavily on what is required by the thinker to whom the explanation is addressed: that thinker whose lack of understanding is an issue. And so if there is such a thing as an explanation *simpliciter* of some phenomenon, in abstraction from the needs of a thinker, then we should expect that explanation *simpliciter* must be specified with a parameter that specifies the kind of general failure of understanding which the explanation *simpliciter* can remedy.

The usual exposition of contrastive explanation, as I said above, proceeds against the background of a concern with the needs of a philosophy of science. One aspect of this is that explanation of *p* is usually taken to involve a specification of (some of) *p*'s causes. We can usefully generalize here. Explanations in general involve specifying that on which the explained phenomenon depends. In some circumstances this means specifying causes, but in the domain of the normative (as argued in chapter 1), it means specifying grounds.

Taking causes rather than grounds to be central, a distinctive problem of explanation seems to emerge. Many natural phenomena have a long causal history: a history specifying what acted when and how that contributed to its current state. For example, a tattered piece of clothing might have a history of being worn on many occasions, each of which contributed in a small way to that piece of clothing's becoming more tattered. And for many natural phenomena (perhaps even most or all) we can specify a causal chain which extends backward from that phenomenon through each of its causes. The cause of a bacterial infection may result from a lack of washing one's hands

properly, which may be caused by the high price of soap, which may be caused by industrial strikes in the soap industry, and so on.

Explanations are selective, however; they typically consist in giving only part of a phenomenon's causal history. Why did the Colonel murder the Marquis? Because the Marquis had slighted him over croquet. This is an adequate explanation, even though the complete causal history of the murder extends far beyond this: to the motive, the means of the murder, the opportunity for committing it, and so on. So one problem of explanation that seems to emerge is what additional conditions there are on a cause for it to count as explanatory.

This is certainly a genuine problem for the explanation of natural phenomena, but some of the assumptions in the set-up of the problem cannot be imported easily into the normative domain. One (correct) assumption is that the causes of natural phenomena will be many and varied far beyond what is needed for the purposes of explanation. It is not obvious, in contrast, that the grounds of normative truths are too many and varied to fit into an average explanation of that truth. So this consideration leaves open whether we are to think of an explanation of a normative truth as something more restricted than giving the grounds for it (I do eventually endorse this proposal, though, but for different reasons).

A second, and more important, assumption is that we can identify the causes independently of engaging in explanation. If a phenomenon's causes were simply those things given in an explanation of that phenomenon, then no sense could be made of the idea that explanation selects from a phenomenon's causal history. But this is an idea the above problem invokes, so it thereby invokes the priority of causes over explanation. Van Frassen endorses this:

‘Could it be that the explanation of a fact or event nevertheless resides solely in [its] causal net, and that *any* way of drawing attention to it explains? The answer is *no*; in the case of causal explanation, the *explanation* consists in drawing attention to certain (‘special’, ‘important’) features of the causal net.’³⁶ (his italics)

³⁶ Scientific Image (1980), p.124-5, section 2.7

There are two distinctions here which we must not confuse. One is an epistemological distinction: causes are known through science, and explanations are known by getting at the relevant causes. We need not accept this distinction: science itself is partly driven by a quest for explanations. Van Frassen does not accept it³⁷. In any case, it does not mark a contrast with the domain of the normative. There is no way in which we discover the grounds for normative truths apart from engaging in the practice of justification and reason-giving, which I have claimed are explicable in terms of the very notion of explanation.

The second distinction is constitutive: the causes are, essentially, ‘whatever structure of relations science describes’³⁸, and the explanations are whichever causes are relevant, where ‘relevance’ itself has no constitutive relation with the causes. This is a putative asymmetrical dependence: the nature of explanation depends on the causes, but the causes do not depend on the nature of explanation. So the causes of a natural phenomenon exist independently of any explanations of it. In the development of my theory, I will reject this assumption for the domain of the normative. The grounds of a normative truth do not exist independently of the justification for it.

Because I reject this assumption, the problem of causal explanation posed above fails to get a grip for me in the domain of the normative. It is not as though there is a whole net of reasons for a normative truth which a justification then has to prune down, in order to get the relevant ones. Instead, the reasons for a normative truth must be discerned in light of the conception of how that normative truth is to be justified. I do not deny that sometimes there are more reasons for a normative truth than we are interested in, and so within a justification we will want to select. But I do deny that the reasons are constitutively downstream: that the reasons for a normative truth are independent of the justification for it. All this is so far just assertion, of course, but part of the point of this section is to point out potential limits to the analogy between explanation of natural phenomena and the explanation of

³⁷ ‘Explanatory power as a criterion of theory choice. That this is indeed a criterion I do not deny.’ P.23, *ibid.*

³⁸ P.124, *ibid.*

normative truth, even though I will make use of views of the former in order to illuminate the nature of the latter.

Suppose we accept the above problem, however: we need a condition which selects a relevant cause from among all the causes of some event or some object's causal history in order to yield an explanation of the event or state of affairs to be explained. This condition must arise from some structural feature of explanation. Let's call this condition a 'condition of relevance', since it is meant to pick out the relevant cause from among all those involved in the causal history of some event or state of affairs.

Sometimes the cause which figures in the explanation varies, depending on the explanation being asked for. If we ask why a steak knife cuts well, the explanation might vary depending on our interests. One explanation might be that the fineness of the edge makes penetration easier. Another might be that steak knives are designed to cut especially well because of the relative toughness of that which they are paradigmatically supposed to cut. These are different causes, not incompatible with each other, but still each may properly figure in an explanation, depending on one's interests when one asks for an explanation. The amateur physicist may prefer the former answer, the potential customer another. Interests, then, sometimes determine the condition of relevance.

Leaving aside the particular relevance of one's interests in a request for explanation, such phenomena bring out an important general aspect of explanation. Explanations answer intrinsically to requests or demands for explanation. What explanation is appropriate depends on what kind of explanation is being asked for. So there is a separation between an explanatory question and an explanatory answer to such a question. We are to account for conditions of relevance by adverting to features of explanatory questions. And leaves in place a separate inquiry concerning what counts as a good answer to a specific why-question.

We have seen how interests may bear on the condition of relevance at play in an explanation. But there is an additional dimension of explanation that is brought out by asking why something happened instead of or rather than a

different thing. Explanation of actions is an especially natural site for such concerns. Why did the Colonel murder the Marquis instead of chastising him? Why did the Colonel murder the Marquis rather than arrange for a duel? The answer to these questions will differ even though they concern the same event, namely the murder. An answer to the first may mention the Colonel's extreme disposition to anger. An answer to the second may mention the Marquis' excellent skills with a rapier.

Both these questions can be seen as a particular specification of the core question: why did the Colonel murder the Marquis? And they bear on the answer to that general question. Two questions may concern the same event even if they are not derived from a common question using the 'rather than' operator. Why are you wearing tattered clothes? – I can't afford to buy new ones. Why are your clothes so tattered? – I've had them for a long time. These are different questions, and different answers, though they both concern the same state of affairs, namely the state of the clothes the person is wearing.

In the above examples, the difference in questions emerges in a difference in what the actual state of affairs is being contrasted with. The Colonel murdered the Marquis, and this is contrasted with a scenario in which he merely chastised him, or alternatively, another scenario in which he arranged a duel. Someone is wearing tattered clothes, and this is contrasted with the scenario in which they are not wearing tattered clothes, and alternatively, another scenario in which those clothes they are wearing aren't tattered.

In each of these examples, then, a contrast scenario is introduced. It is contrasted with what is actually the case. Following van Frassen³⁹, we call the actual state of affairs the 'topic' of a request for explanation. In the above examples an adequate explanation must explain why the topic is true rather than the contrast. Why did the Colonel murder the Marquis rather than chastise him? The Colonel's extreme anger explains why. But the anger wouldn't explain why the Colonel murdered the Marquis rather than arranging a duel; both would be adequate expressions of anger. So we must introduce another explanatory factor, the unlikeliness of a successful duel.

³⁹ P.141, *ibid.*

This explains why the Colonel would choose to murder the Marquis rather than fight him honourably. It is tempting to try to analyse away the contrast and replace it with a non-contrastive explanandum: perhaps a conjunctive one, such as ‘Why did the Colonel murder the Marquis and not arrange for a duel?’ But this is misguided; I defer to Lipton’s excellent arguments here⁴⁰.

It is important that the contrast is not always incompatible with the topic. Although one can ask why the Colonel murdered the Marquis rather than chastising him, it is not ruled out by the question that he could have done both. Sometimes it is incompatible, as when one asks why someone’s clothes are tattered instead of whole: they couldn’t have been both. Moreover, it is important that the contrast does not always concern the same event as the topic. One can ask why the Colonel murdered the Marquis but he did not murder the Earl, even though the Earl also slighted him over croquet, just as the Marquis did. This question does not ask of one event why it did not turn out differently, but why a different event did not turn out the same, as we might have expected it to.

All this raises the question of what conditions there are on something to count as a sensical contrast with a topic. It doesn’t have to be incompatible, and it doesn’t even have to concern the same event or state of affairs. But there are clearly limits. It doesn’t make sense to ask why the Colonel murdered the Marquis but he did not wash up the dishes. But is there an illuminating abstract condition on something to count as a contrast? The closest that can perhaps be said is that something counts as a contrast relative to a given topic if both involve the same kind of causes, and there is a question over how those causes operate. For example, the Colonel murdering the Marquis involves the same kind of cause as chastising him would: both would be a response to the slight over croquet. The question is why the slight over croquet resulted in the topic rather than the contrast. On the other hand, if we ask why the Colonel murdered the Marquis but did not murder the Earl, even though both slighted the Colonel over croquet, again we have the same cause, namely the croquet insult. There is a question over what effect that insult had: why it had one

⁴⁰ See the section entitled ‘Failed reductions and False differences’, in chapter 3 of *Inference to the Best Explanation*, 2nd ed. (2004)

effect in the topic but not the same effect in the contrast. No common cause is present in the dishwashing contrast: murdering the Marquis, and washing up the dishes, would not share a common cause. Showing that there is no sensical contrast is a way of showing that a request for explanation which uses a particular contrast is not a legitimate explanation request.

Some, perhaps all, of the relevance of one's interests to explanation can be captured in the selection of a suitable contrast or class of contrasts. The amateur physicist wants to know why steak knives cut well while some other objects which physically differ don't. The potential customer wants to know why steak knives cut especially well while various other consumer goods don't. Instead of speaking of explanation as relative to interests, then, we should say that an explanation of a topic is an explanation relative to a particular class of contrasts (and the selection of contrast class may depend on one's interests). When we vary the contrast class, the explanation varies. Because of this, giving an explanation of a topic never requires the selection of a particular contrast class instead of others. The explanation itself pertains only relative to a given contrast selection; and as long as the contrast class one selects is genuine, there are no constraints on which contrast class one must choose.

This raises the question of whether or how it is possible to explain a topic *tout court*, absent any particular contrast class. We cannot take this for granted on the basis that ordinary-language requests for explanation infrequently specify an explicit contrast class, for a contrast class may frequently be implicit. For example, if someone asks why the Colonel murdered the Marquis, period, then we might suspect that they are asking why the Colonel murdered the Marquis as opposed to not wanting to kill him at all, a question answered by a mention of the insult over croquet. That would certainly be a reasonable interpretation of the question.

Lipton has another suggestion: he assimilates explaining a topic *tout court* to explaining why it should happen instead of its being the case that it does not happen. He says:

‘The way to construe [our preferred conception of explanation] as it applies to the limiting case of the contrast, P rather than not-P, is that we must find a difference for events logically or causally incompatible with P, not for a single event, ‘not-P’... To explain why P rather than not-P we do not, however, need to explain every incompatible contrast. We do not, for example, need to explain why Jones contracted paresis rather than being long dead or never born. The most we can require is that we exclude all incompatible foils with histories similar to the histories of the fact.’⁴¹

The suggestion is in effect that we quantify over contrasts and explain why *p* as opposed to many incompatible contrasts i.e. many ways in which not-*p* might have occurred. The suggestion is correct in that explaining all of these incompatible contrasts will be relevant to giving a full explanation of *p* (I assume that explaining *p tout court* is the same as explaining why *p* as opposed to not-*p*). Explaining why the Colonel murdered the Marquis involves explaining why he did so rather than chastise him or walk away. Is Lipton right to suggest that there are limitations on ways of being not-P that are required in order to explain why P rather than not-P? Although it seems plausible that explaining why Jones contracted paresis shouldn’t require explaining why he isn’t long dead, one wonders why. The reason, I think, is that this way of being not-P is not a contrast to P in the first place, even though it is a way of being not-P. It wouldn’t make sense in the first place to ask why Jones contracted paresis rather than being long dead. Hence we should think of Lipton’s suggestion of universally quantifying over contrasts: as using the criterion for something’s being a contrast in order to determine what ways of being not-P require explanation when trying to explain why P rather than not-P. So, in this revised proposal, giving an explanation as to why *P tout court* is explaining why P rather than not-P, which in turn is explaining why P as opposed to every way in which not-P might have been the case.

Another question we might ask, though, is why explaining *p* rather than not-*p* should be limited to incompatible contrasts. Why should explaining why the Colonel murdered the Marquis instead of leaving him alone not also involve explaining why the Colonel murdered the Marquis but not the Earl? Explaining this latter contrast may shed light on the Colonel’s motives for

⁴¹ p.49-50, *ibid.*

murdering the Marquis, and as such are relevant to the question of why he murdered the Marquis, *tout court*, and are part of a full explanation of why he murdered the Marquis. So for example, if it turned out that the reason the Colonel murdered the Marquis but not the Earl was because he was friends with the Earl and he thought that he should forgive the Earl out of friendship, then part of the explanation for why he murdered the Marquis is that the Colonel did not have a reason of friendship to forgive the Marquis' insult.

Finally, a couple of remarks on what makes a legitimate answer to a (legitimate) request for explanation. The key is that an answer discriminates between the topic and the contrast: it provides a difference between the two that explains why they turned out differently, with the implication that had the difference not been in place, the two would have turned out the same (if there are no additional relevant differences either). For example, if the reason the Colonel murdered the Marquis but not the Earl, even though both insulted him, was that he had great respect for the Earl but not for the Marquis, then had the Colonel had great respect for the Marquis as well then he would not have murdered him (if there are no other relevant differences); conversely, if the Colonel had had no respect for the Earl, he would have murdered him (if there are no other relevant differences).

A tricky issue for theories of explanation that take their cue from the philosophy of science is how to construe explanations that deal with non-determining causes i.e. causes that make an outcome probable rather than certain. For example, if being an atom with a certain unstable internal configuration makes it 50% likely that it emits a particle within a minute, and it does, how good an explanation of this emission is it that the atom had that unstable configuration? This issue does not arise for norms, however, given the assumption that there are no normative truths whose justification make them merely probable. There may be normative truths which are themselves probabilistic truths, as when it is probable that one course of action in particular is the best. But this is not a truth whose justification makes it merely probable: the justification may be perfectly adequate. This is another way, then, in which a theory of explanation for normative truth differs from that for scientific truth.

Now that I have laid out a certain conception of explanation, I will illustrate how this conception can be applied to the explanation of normative truth. First I will discuss justification for normative truth: what justification is if it is an explanation, and if explanation works contrastively. Then, in the next chapter, I will turn to the topic of reasons, and how reasons figure in justifications: I conceive of this topic as very much downstream from the nature of justifications.

Justifications are explanations, and explanations are contrastive: in every explanation there is a parameter, either implicit or explicit, which is the set of contrast cases. So we should expect the same to be true of justifications. Within every justification, there is a parameter which is the contrast class. Every scenario in the contrast class contrasts with the topic being justified, and each provokes a burden of justification, in the same way that for explanation a topic-contrast pair provokes a burden of explanation. Here are some examples of normative topics and possible contrasts.

- Choosing to undergo plastic surgery is in general a dodgy thing to do, if not actually wrong. While there is nothing wrong with reconstructive surgery in particular, there is something amiss with undergoing plastic surgery in order to try to make oneself more attractive. However, there is nothing wrong with putting on makeup. It is perfectly legitimate to put on makeup to try to make oneself more attractive, within limits—people should avoid damaging their skin, or wearing makeup compulsively or out of insecurity. Both plastic surgery and makeup involve changing one’s appearance to make oneself more attractive. The question is unavoidable: Why should it be that choosing to undergo non-reconstructive plastic surgery is untoward, while wearing makeup is in general OK?
- The final round of the 2002 French presidential election featured two candidates: Jacques Chirac, and Jean-Marie Le Pen. It was obligatory to vote for Chirac in this election. But it wasn’t obligatory, or at least not as obligatory, to vote for Chirac in any of the previous elections. There are two questions. Why was it obligatory to vote for Chirac over Le Pen? (The answer should be known to anyone with a minimal

knowledge of the two candidates.) Second: Why the special urgency in the 2002 presidential election?

- Harry faces a choice over what kind of career to pursue. One option that is a real possibility for him is to become a well-paid IT consultant: he's been offered a job of this nature. A second possibility for him is to become an entrepreneur and pursue his own ideas, of which Harry has many, but with all the usual risks: low income, failure. So the question emerges for Harry: how shall he live? Suppose he should become an entrepreneur. Why should he become an entrepreneur instead of a consultant?

These examples differ in form.

The first contains a topic to be justified/explained, namely the dodginess of plastic surgery, and the contrast is the non-dodginess of makeup. It is striking that the one should be dodgy and the other not, given their seeming essential similarity. For anyone who believes that plastic surgery is dodgy but that makeup is not, the question arises as to why. This question also arises for anyone who is *tempted* to believe that plastic surgery is dodgy but who is not tempted to believe that wearing makeup is. Because of these features, it is a question about the justification of the proposition that plastic surgery is dodgy. If no answer can ultimately be given to this question, then it cannot be true that plastic surgery is dodgy yet makeup not. If it were true, it would have an answer. So the burden of justification does not merely arise for someone who already believes the two claims. It also arises for anyone considering these two claims: unless the question is settled for them, they cannot consider the proposition that plastic surgery is dodgy justified.

In this example, the contrast is compatible with the topic in the sense given above that it is not an alternative to the scenario envisaged in the topic. It could be true both that plastic surgery is dodgy and makeup is. It could also be true both that plastic surgery is not dodgy and wearing makeup isn't either. Finally, it could be true that the one was dodgy and the other not, and vice versa. All these are possible positions in logical space. In addition, there is a clear common cause or common ground that might potentially explain why

there is a pressing contrast here. Both procedures involve changing one's appearance in order to make oneself more attractive. Given this particular similarity, it is surprising that the two procedures should ethically differ. In this way the example exhibits some of the hallmarks of a contrastive burden of explanation explored above.

In the second example, there are two separate contrastive burdens of explanation/justification. First, there is a burden of justification concerning why it was obligatory to vote for Chirac over Le Pen. Unlike the previous example, this contrast is incompatible with the topic, in that the scenario it envisages forms an alternative to that envisaged in the topic. The choice at the time was to vote either for Chirac or for Le Pen, and voting for both was impossible. In this example there is no obvious common ground, no relevant similarity which might explain why the contrast is there. For this example, the explanation of why there is a contrast there is bound up with the importance of the options. There was much at stake in that election, and therefore it was of importance who to vote for. Hence, if it was obligatory to vote for Chirac over Le Pen, a pressing question arises over why the one rather than the other.

The second burden of justification is quite different, concerning why it was obligatory to vote for Chirac in the second round of the French 2002 election, but not obligatory, or at least not quite as obligatory, to do so in any of the elections prior to that (or any of the ones after) in which he stood. This contrast is compatible with the topic: it could be true that it was obligatory to vote for Chirac in all these elections, or in none of them, or in any particular set of them. Furthermore, there is a common ground: the fact that Chirac was on the ballot in each of them.

From these examples of compatible contrasts and their common grounds I tentatively conclude that the existence of the common ground (and hence the existence of the contrast) is related to the reasons in favour of the topic, and depends on them. Part of what made it obligatory to vote for Chirac over Le Pen in that 2002 election was certain features of Chirac that Le Pen lacked. Hence, the fact that it was Chirac standing is part of the reason for voting for

him over Le Pen in that election. The fact that it was Chirac standing was the common ground we adduced to explain the existence of the contrast class composed of other elections where Chirac stood. A question emerged over why it was obligatory to vote for Chirac in that 2002 election but not in the others. Likewise, with the example of plastic surgery, that it is done solely to make oneself more (conventionally) attractive is arguably part of the reason for its dodginess. Perhaps it represents an uncomfortable subversion of a hallowed personal space before the external authority of social convention. If this is true, then the fact that it is done in order to make oneself more conventionally attractive should serve as a common ground that creates a contrast wherever it pertains in a different circumstance. It pertains in the case of makeup, hence it creates the contrast of makeup to the topic of non-reconstructive plastic surgery. It creates the question concerning why makeup should be legitimate and non-reconstructive plastic surgery not.

There is only one contrast in the third example, and it is an incompatible contrast: Harry cannot be both a consultant and an entrepreneur. Much like the first contrast in the second example, what makes it a contrast is the importance of the choice: it will bear heavily on his future choices and, if he makes the wrong choice, it will bear on his quality of life. Hence, if Harry is tempted to choose one over the other, there is a pressing question for him as to why he should choose the one over the other, if he should do this at all. If an answer cannot be given, then the proposition that he should the one over the other is unjustified.

The purpose of these examples has been to illustrate how contrastive burdens of explanation apply with full force to justification, to make it clear how there can be contrastive burdens of justification as well. Another purpose of these examples is to illustrate the generality of the thesis. A similar kind of demand for justification applies in the case with an overtly moral claim about who it was obligatory to vote for; one concerns principally a question of prudential evaluation; and finally, the first example concerning plastic surgery is in the realm of the ethical without being overtly moral.

What is of interest in these examples is the kind and origin of the questions that arise; I will not look at specific answers to these questions, though I strongly expect that there are adequate answers. If there are answers, then they are answers that provide explanations. If we understand why plastic surgery should be dodgy but makeup not, then we know part of the answer as to why plastic surgery should be dodgy, and so we know part of the justification as to why plastic surgery should be dodgy. Suppose the answer is that plastic surgery makes one to some extent unrecognizable, whereas makeup generally doesn't do this. Then we know at least part of why plastic surgery is dodgy—because it makes one unrecognizable—and we know at least part of the justification for its being dodgy—because it makes one unrecognizable. So the answers to such question are explanations and thereby justifications, and the questions that elicit them are therefore themselves requests for justification and thereby requests for explanation.

I have said that the presence of a common ground is a way a contrast can come to be (in the examples so far, one which is compatible with the topic). What is common must be a ground in the sense that it must be closely related to a reason in favour of the topic. For example, the fact that Chirac was standing is closely related to the justification for voting for him over Le Pen in the 2002 second-round French presidential election, because the fact that Chirac was standing is presupposed by the justification for voting for him over Le Pen. Hence, where this pertains elsewhere, the potential for a contrast is created. If it was obligatory to vote for Chirac in the second round of the 2002 election but not in other elections where he stood, the fact that it was Chirac standing is common throughout and so a serious question arises as to why it was obligatory to vote for him in that election but not in those other elections.

Where a common ground underlies a burden of justification, one can show the burden of justification to be illegitimate by showing that there is no common ground; if there is a burden of justification, it has to be given some other foundation. The existence of a specific common ground may be a presupposition someone makes if they think that a specific scenario constitutes a contrast to the topic, one which poses a serious question as to

how the topic can be justified. So by showing the presupposition to be false, one can show the question to be illegitimate.

Here is an example. Suppose a modern-day Christian has the following pair of attitudes. They condemn adultery, but they do not condemn divorce. For us the question arises over why they should do one but not the other. After all, the biblical prohibition against divorce is at least as strong as that against adultery. However, this person, when pressed, denies that there is a serious biblical prohibition against divorce. This refutes the question: it no longer makes sense to ask why they should condemn adultery but not divorce, with this knowledge that they deny the existence of a biblical prohibition against the latter.

Common grounds are a source of compatible contrasts, (but not, I will argue later, a necessary condition on their existence). In the case of incompatible contrasts, I claimed that the importance of one's options is a common source in cases of choosing what to do. If Harry ought to choose to become an entrepreneur over a consultant, then this requires a justification, because of the potentially huge ramifications of the choice; and whatever he is uniquely justified in doing (if anything), that is something Harry in some sense ought to do. Even if the two are on a par or are in some sense incommensurable with each other, then this in turn still requires a justification, just because of the importance of the choice.

In contrast, a choice scenario in which there were no potentially important ramifications to the choice would not demand justification in these ways. Say, for example, you have decided to snack on a piece of fruit and you are faced with a choice between some strawberries and some raspberries. Certainly, if it's true that you *ought* to choose the strawberries over the raspberries, then the question arises as to why this should be so. So even in this circumstance of unimportant choice, contrastive burdens of explanation can arise. But if it's not true that you ought to choose one over the other, as is plausible in the case of strawberries and raspberries, then no topic/contrast pair is generated, and no demand for justification arises. This contrasts with the case of important

choice: there, if the two options are on a par, the demand for justification still holds.

Anscombe was right when she wrote:

‘when I do action A for reasons R, it is not necessary or even usual for me to have any special reason for doing-action-A-rather-than-action-B, which may also be possible.’⁴²

She is right in the sense that a requirement to have an adequate justification for doing one thing rather than another does not usually pertain to cases of choosing to do one thing over another. I need have no special reason for choosing outfit 1 over outfit 2 when I choose to wear outfit 1 today. But this feature marks out a contrast between the justification for actions and the justification for normative propositions. Choosing to do one thing over another does not automatically require a justification, but if you *ought* to do one thing over another, this does automatically require a justification.

Because of the closeness of ‘ought to do X’ propositions to actions, it is tempting to assimilate the two vis-à-vis when justification is required. I have argued that this is mistaken. ‘Ought to do X’ propositions are one kind of normative proposition that involves incompatible contrasts and so fall under that more general account. Another kind of normative proposition that involves an incompatible contrast concerns what is optimal. For example, one might wonder whether things go best when the population is pious, or whether things go best when the population is not at all pious. Perhaps things go best when there are varying levels of piety in the population, so there would be diversity in ways of living. This is not a question over what to do, though it might be important to matters of public policy. Still, it involves incompatible alternatives that are interesting i.e. which are not simply the contradictories of whatever proposition we are considering. What burdens of justification are there here, and what generates them?

At this point we should revise the idea that the importance of the issue is what generates the contrastive burden of justification where incompatible contrasts pertain. We can extend the idea that a common ground is what generates the

⁴² Who is Wronged?, Anscombe (1967), pp.16-17

burden, but connect this to the importance of the issue. Suppose that things are optimal when the population has varying levels of piety. This is an issue of importance. Part of the explanation of why things are optimal when the population is variable in piety is that the issue of the general piety is an important one: it affects the ways of life that are open to the average citizen, and affects more besides. If it didn't matter whether the population was pious or not, then any particular configuration of the general piety wouldn't be better than any other; and so there wouldn't be a fact of the matter as to whether things go best when the population is pious or not.

So, the importance of an issue can be part of the explanation of why a particular normative verdict on that issue is appropriate. If the issue weren't important, then no strong normative verdict would be appropriate. Because it is part of the explanation, it is apt to serve as a common ground in the sense explored above. It can be shared between different scenarios and so a contrastive burden of justification is generated: one of those scenarios will be a contrast and another the topic. For example, suppose again that things are optimal when the population has varying levels of piety. The question arises: why is it better that the population is variably pious than that it is totally pious; and why is better that a population is variably pious than that it is not at all pious? Here the contrasts (high piety and no piety) are incompatible with the topic of varying piety. So here the presence of a common ground can make an incompatible alternative a contrast to the topic, generating a burden of justification.

This is the general account that we should apply to the 'ought to do X' propositions as well. If Harry ought to become an entrepreneur, then part of the explanation for this is that it is an issue of importance, affecting what choices he will be able to make in the future. Becoming a well-paid consultant is an issue of importance as well; again it will greatly affect the kind of life Harry will be able to lead. So the question arises: why ought Harry to become an entrepreneur instead of its being the case that he ought to become a consultant?

We have, then, an entirely general account of when contrastive burdens of justification arise: they arise where there is a common ground between contrast and topic, a common ground which would form part of the justification/explanation for each. In cases of incompatible contrasts, this common ground is the importance of the issue. In cases of compatible contrasts, the common ground varies between specific cases. If all burdens of justification are contrastive, then we have a serious constraint on when a need for justification arises. I have not yet argued that all burdens of justification are actually contrastive in nature, even when they do not explicitly mention a contrast. First, I will argue that *if* this is true, a very interesting consequence follows with respect to the possibility of basic normative truths, truths which are not explained by any other normative truths.

Suppose that the existence of a common ground is necessary and sufficient for the generation of a need for justification in the ways I have outlined: we can put aside the issue I raise later about burdens of justification theorised into existence. It follows that where a common ground is absent, no need for justification arises. Further, it follows that if no common ground exists for a particular normative topic with respect to any possible contrast, then there are no contrasts for that topic. That topic would then lack a need for justification. Furthermore, there would be no respect in which a justification could be provided: for that a justification would have to address contrastive burdens of justification which, *ex hypothesi*, do not arise.

Let's call this property of not needing a justification the *terminal property*. A normative proposition with the terminal property lacks a need for justification and lacks the possibility of being justified. This is not to say that that proposition must on that account be true. Rather, it is to say that if it is true, still it does not need justification, nor could one be provided. If it should be given up, this requires a different ground than not being justified.

In Chapter 1 I assumed an understanding of justification according to which if a justification could not ultimately be provided for a proposition that needs one, then we shouldn't take that proposition to be true. This does not imply that any proposition which ultimately lacks a justification should be given up,

because some propositions may not need justification. A non-normative example of such a proposition is perhaps the law of non-contradiction: a proposition and its contradictory cannot be simultaneously true. If it is true, then it does not seem as though it is true in virtue of some other truth, or that it is to be explained by some other truths. Still, it does not seem to hold brutally: it is not a cost to our view of the world that we cannot give a deeper explanation of the law of non-contradiction. In the same way, if some normative propositions are terminal, it may not be a cost to our normative worldview that we cannot give a deeper explanation or justification for them.

What kind of normative proposition is a candidate for being terminal? I suggest an Aristotelian-in-spirit function principle as an illustrative example. I am not concerned with the function principle qua argument, so that we can deduce rich ethical conclusions by filling information about what the functions of humans are. Furthermore, I am taking the function principle out of the strong naturalism in which it is embedded, and treating it as a self-standing principle. Hence the principle is Aristotelian only in spirit. Let's call this principle G:

G): It is good when humans act so as to develop their distinctively human capacities

To show that G is terminal, we have to show that there are no potential contrasts that could be drawn.

Are there any compatible contrasts? Let's try a few. It is not good (it need not be bad, but still is not a good thing) when humans act so as to develop distinctively piggish capacities—let's assume this. (Let's call this principle H). So perhaps a question arises concerning why it should be good when humans develop distinctively human capacities but not good for them to develop distinctively piggish capacities. This is compatible, since it is a possible position in logical space that both are true: it is good when humans develop distinctively human capacities and not good for them to develop distinctively piggish capacities.

The common ground would be distinctive capacities: since its being distinctive capacities that are at stake is important to G, and they are also at

stake in H, a question arises over why the two should differ. But the common ground is not genuine. The presence of distinctive capacities does not explain G in any way or extent, only the distinctively *human* capacities. Hence there is no common ground between G and H on this score, and there is no burden of justification as to how G and H can both be true.

What about: it is good for humans to develop distinctively human capacities, but not good for pigs to develop distinctively human capacities. Why? But in this example, it is not a real possibility that pigs *could* develop distinctively human capacities. Hence, it is questionable whether there is a genuine burden of justification here.

In the roughly Aristotelian framework which I am considering, it is good when any x behaves so as to develop or fulfil its distinctive capacities, or those capacities distinctive of x's kind, not just when humans do so. So no contrast can be generated on that score.

When a normative truth is taken as basic, it is inevitable within my system that there will be no compatible contrasts, even if the truth ultimately isn't basic, *contra* how it is taken. For a compatible contrast requires a common ground, a common ground is part of the justification for those truths, a truth which is taken to be basic is taken not to have a justification, and so any normative truth taken to be basic must be taken not to have anything which could serve as a common ground and generate a contrast.

What about incompatible contrasts? The simplest incompatible contrast is the negation of G: it is not good when people act as to develop their distinctively human capacities. Does this create a burden of justification? Only if there is a common ground resulting from the importance of the issue.

Cases of incompatible contrasts created by the issue's being important were like the following. It is important if it is optimal that the population is pious, since it has a significant effect on proper policy decisions. Similarly, it is important if it is optimal that the population is not at all pious: that likewise has a significant effect on proper policy decisions.

This schema, interestingly, cannot apply to G. It is important if G is true, since that affects our evaluation of all kinds of things. But it is not important if G is false, since if we start out from a position of neutrality on G, the conclusion that G is false does not affect our evaluations. If we were neutral on G, we were not guided by G in making our evaluations, and so deciding against G does not change any of the evaluations we made. The only sense in which it is important that G is false is the derivative and uninteresting sense that it is important if we were mistaken in thinking G to be true. So the importance of an issue cannot be the common ground which generates an incompatible contrast. Hence, there is no burden of justification on why it should be true that it is good when humans develop their distinctive capacities, as opposed to false.

Here is another incompatible contrast: it is bad when humans develop distinctively human capacities. We might wonder: why should it be a good thing when humans develop distinctively human capacities as opposed to a bad thing? Both of these would be important: if it is a bad thing that humans develop distinctively human capacities, then this should affect everything from policy decisions to personal life decisions. So it seems that an incompatible contrast is generated here.

But the importance of the issue does not genuinely generate a contrast here. For it to do so, it would have to be part of the explanation of why it was good that humans develop their distinctively human capacities. But that that issue is important is not part of the explanation as to why human development of human capacities is good. Rather, it's being good explains why it is important. In contrast, in the case of career choice, that career choice is an important issue is part of why one ought to make a particular choice over another: it's not the case that the fact that one ought to choose a particular career over another is part of the explanation as to why career choice is important. (Incidentally, this suggests that in general 'ought to do X' propositions are more in need of justification than propositions concerning what is good or bad.) So, G appears to be genuinely terminal.

This discussion illustrates how a normative proposition can fail to incur a contrastive burden of justification. On the assumption that all burdens of justification are contrastive in nature (I will argue for this soon), that normative proposition is terminal: it incurs no burdens of justification *simpliciter*. Since justifications essentially respond to burdens of justification, it follows that for a terminal normative proposition there is simply no justification. So it is basic.

This is not the only way in which a particular normative truth might be thought to be basic. Instead, some particular normative truth might be posited by theory to be basic, even though it clearly faces a burden of justification. For example, suppose we believed that it is basic that it is impermissible to kill. More likely we would simply not be articulate enough to be able to give a better justification. Still, it is possible that it is ‘just’ impermissible to kill—without there being a deeper justification for this, a justification which may fail to apply in some cases. In this theory, then, it is basic that it is impermissible to kill, even though, there is a clear burden of justification as to why. Because of this, it is a bad theory. A better theory would give a deeper justification. Still, this shows how it can be coherent to posit a particular normative truth as basic, as lacking a deeper justification, even when it is in need of one.

There is an analogy for this in science: it is possible for a theory to take something as basic, lacking a need for explanation. Van Fraassen gives several illuminating examples:

‘We also reject such questions as the Aristotelians asked the Galileans: why does a body free of impressed forces retain its velocity?...Clark Maxwell accepted as legitimate the request to explain electromagnetic phenomena within mechanics. As his theory became more and more accepted, scientists ceased to see the lack of this as a shortcoming. The same had happened with Newton’s theory of gravitation...In both cases there came a stage at which such problems were classed as intrinsically illegitimate, and regarded exactly as the request for an explanation of why a body retains its velocity in the absence of impressed forces...the important fact for the theory of explanation

is that not everything in a theory's domain is a legitimate topic for why-questions: and that what is, is not determinable *a priori*.⁴³

It is possible, then, to theorise away a burden of explanation. But likewise, it is possible to theorise them into existence. An interesting example within the normative realm, of introducing a burden of justification, arises from the demands of freedom. Rawls claims:

‘A...respect in which citizens view themselves as free is that they regard themselves as self-authenticating sources of valid claims. That is, they regard themselves as being entitled to make claims on their institutions so as to advance their conception of the good.’⁴⁴

For Rawls, if citizens make claims on institutions to advance their conception of the good, and those institutions refuse, a burden of justification arises on institutions to be able to justify this refusal: why may they refuse rather than not? For one cannot gratuitously throw out valid claims. So this is an example of a theory positing a special burden of justification that goes beyond the common ground present in generic contrastive burdens of explanation. In this way, it is a mistake to think that a common ground is a necessary condition for there to be a contrastive burden of justifications: it is a normal condition, but nothing rules out our positing another source of contrastive burdens of justification.

We have now explored the origins of contrastive burdens of justification and seen to when they may also fail to hold. However, we have not yet addressed the question of how far this account can be extended. While it is clear that there are genuine contrastive burdens of justification, it is not yet clear that all genuine burdens of justification are contrastive in nature. There are two kinds of cases in particular that are of interest. The first concerns implicit contrasts: cases where a burden of justification is posed without an explicit contrast, but where a contrast must be recovered from the context and information about the likely interests of the inquirer. The second concerns global contrasts: cases where we want to know if overall there is an adequate

⁴³ P.111-2, *ibid*.

⁴⁴ Rawls (2005) Political Liberalism, part 1, lecture 1, part 5: The political conception of the person, section 3.

justification for a topic. Both of these cases are legitimate questions, and the task facing my account of normative justification is to construe these legitimate questions in terms of other legitimate questions that are acceptable to the contrastive theory of justification.

As the name suggests, the key to dealing with cases of implicit contrast is to impute a hidden contrastive structure to the explicit burden of justification. Here is an example of a burden of justification with only an implicit contrast: it is election time and I am wondering if I should vote for a particular candidate—let's call that candidate Dean. This is a burden of justification on whether I should vote for Dean with no explicit contrast. An implicit contrast can be imputed easily enough: what I am wondering is, for each other candidate, why it should be that I should vote for Dean instead of its being the case that I should vote for that other candidate. Finally, what requires justifying is that I should vote for Dean instead of not voting at all, or spoiling my ballot. The burden of justification as to whether I should vote for Dean is equivalent to all of these other ones put together. I suspect that the lesson here can be applied fairly generally: where the contrast must be implicit if present at all, the implicit contrasts tend to be incompatible and wide-ranging.

Global contrasts are much more interesting, where the burden of justification for a topic is unrestricted: what we are wondering is whether the topic can be fully justified. Global contrasts too can be implicit, and sometimes it is difficult to discern whether a burden of justification presented without an explicit contrast has an implicit global contrast, or a more restricted implicit contrast. This is not a philosophical question so I shall not pursue it further. It does bring out an important point, though: there is no ordinary language notion of a global contrast. A question arises over whether I should vote for Dean, but that can be interpreted as a non-global contrast as I did above. So even in asking whether a particular topic is fully justified, or justified *simpliciter*, we are already in philosophical space.

This means that we have some license to define the notion of being fully justified as we see fit. It follows from what I have already said that it is a necessary condition on a normative proposition's being fully justified that it

be able to meet a contrastive burden of justification. If no response can ultimately be given to a burden of justification, then the topic under that burden is unjustified. So, for example, if no response can ultimately be given to the question of why plastic surgery should be dodgy but makeup not, then the claim that plastic surgery is dodgy is not fully justified (though one may have plenty of other good reasons for thinking it).

Another question is whether it is sufficient for a normative proposition's being true that its responses can ultimately be given to all of its contrastive burdens of explanation. If this is false, then there is an aspect of being justified that cannot be captured contrastively, and that would be a lacuna in my account. I argue that there is no aspect of being justified that cannot be captured contrastively, and so I argue that meeting contrastive burdens of justification is sufficient for a normative proposition to be justified. But this requires an additional proof which I will now provide.

The proof will proceed via *reductio*. I assume, I hope uncontroversially, that being fully justified is necessary and sufficient for a normative proposition to be true. Hence, the outcome of the proof is that iff adequate responses can ultimately be given to all of a normative proposition's contrastive burdens of justification, that proposition is true.

Suppose this condition wasn't sufficient to be fully justified. Then, there would be a normative proposition *P* such that all the contrastive burdens of justification its truth introduces could be adequately met, and yet *P* was still less than fully justified, and therefore false. As long as it is non-terminal, it could be true only if it could be fully justified.

Since *P* is false, by logical equivalence, the negation of *P* would be true. By the necessary condition on full justification discussed above, all the contrastive burdens of justification introduced by the truth of $\neg P$, if there are any, could in turn be adequately met. All the contrastive burdens of justification introduced by $\neg P$ could be met, but *ex hypothesi*, all the contrastive burdens of justification the truth of *P* introduces could also be adequately met.

Whenever a contrastive burden of explanation occurs and the contrast is compatible with the topic, the explanation of the compatible contrast will be at least part of the explanation of the incompatible contrast as to why the event happened rather than not, or why it was the case rather than not. All answers to contrastive questions such as ‘why did the Colonel murder the Marquis but not the Earl?’ will bear on the answer to the core question ‘why did the Colonel murder the Marquis?’. Another example: It is appropriate for the General to assault Fort Blue but inappropriate to assault Fort Red, though it is in his power to do both. This is because there are high-level targets in Blue, but none in Red. This explanation of a compatible contrast will also serve as part of the explanation for the topic (it being appropriate to assault Fort Blue) over its contradictory. Why it is appropriate to assault Fort Blue instead of its not being appropriate to do so is partly explained by the presence of high-level targets in Fort Blue.

It follows from this premise that where there are any contrastive burdens of explanation for a normative proposition, there is a burden of explanation for that proposition over its contradictory. It follows, then, that for P above (which by assumption isn’t terminal) there is a burden of explanation arising as to why P rather than \neg P. Likewise, there is a burden of explanation arising as to why \neg P rather than P. *Ex hypothesi*, all the contrastive burdens of explanation the truth of P introduces could also be adequately met, and the same for \neg P. Since P vs. \neg P is one of these contrastive pairs, as is \neg P vs. P, it follows that there is an adequate explanation for why P rather than \neg P, and *also*, an adequate explanation as to why \neg P rather than P.

This is a contradiction. We are assuming that explanation of normative propositions does not proceed by probabilities. If there were, then it need not be a contradiction. For example, suppose that we are trying to explain why an unstable atom decays after more than one minute: it has a 50% chance of doing so. Why should it decay before one minute? Again, it has a 50% chance of doing so. If both of these are reckoned to be adequate explanations, then this would suffice as a case where we could adequately explain both why P rather than \neg P, and also why \neg P rather than P. But since I assumed in chapter 1 that normative truths do not ever work like this, then we can conclude that

it is impossible for there to be both an adequate explanation as to why P rather than $\neg P$, and also as to why $\neg P$ rather than P . Because it is a contradiction, the original assumption must be false. Therefore, iff adequate responses can ultimately be given to all of a normative proposition's contrastive burdens of justification, that proposition is fully justified (which is sufficient for its truth).

This completes my discussion of burdens of justification/explanation for normative propositions. All burdens of justification/explanation for normative propositions are contrastive in nature. The existence of contrasts to a given topic seems to require a condition that there is a common ground between it and the topic, where the common ground is part of what explains the truth of the topic. This allows for a certain kind of basic normative truth, over and above that kind of basic truth which is posited by theory. In the next chapter I will look at what answers are to this kind of burden of justification, and in particular how this suggests we should conceive of reasons and their relations to each other.

Contrastive Explanation in the Structure of Justification

In this chapter I will pursue further the idea of justification as contrastive explanation outlined in the last chapter. In this chapter I will explore the nature of reasons, but reasons conceived in the light of this conception of justification. I will argue that we should see reasons themselves as a kind of explanation, as those things which can lift the contrastive burdens of justification.

Sometimes a single reason can by itself be a completely adequate justification. We shouldn't jump off the Eiffel Tower without a parachute for the single reason that it will kill us. That single reason is all that's needed. Sometimes an adequate justification is formed by multiple reasons: we should give to charity because by doing so we help victims, and also because we develop our own benevolent sensibilities by doing so. Reasons, then, are sometimes justifications, and sometimes are parts of or elements in a justification. I take this as a datum, and I aim to account for it.

In the last chapter I argued that a burden of justification always featured a contrast class as a parameter. So we can conceive of a justification as whatever answers to a particular burden of justification. Putting this together with the datum expressed in the paragraph above, it follows that sometimes a burden of justification can be answered by a single reason, and sometimes a burden of justification can be answered by a justification composed of multiple reasons.

An example of the former is the following. Take the burden of justification: why is exercising a good thing, rather than morally indifferent? An adequate justification for this is that exercising greatly improves one's life expectancy. That it improves it is a reason for exercise's being good. The reason is the justification, which is that which answers the contrastive burden of justification.

An example of the latter is the following. Take the burden of justification: why is it wrong, rather than morally indifferent, to waste a government

student subsidy not doing any work? An adequate justification for this, in my opinion, is that it is objectionable on a personal and political level: personal because doing so neglect's one's own talents, and political because doing so undermines the general political will to keep such subsidies going when this will is worth preserving. Here the justification is that which answers the contrastive burden of justification, but the justification is composed of multiple reasons.

In some kinds of multiple-reason cases, the justification is overdetermined. For example, it is wrong to gratuitously start a fire in someone's house because it will both destroy their home and also waste the attention of the authorities who might be saving someone else's home in the meantime. Both of these is an adequate justification on its own for the wrongness of the act: even if the other was not an issue, it would still be wrong to do it as long as one of them was in place. We can think of this as a case where there are multiple justifications for the same truth.

Other multiple-reasons cases are most perspicuously construed as a combination of distinct goods. For example, it is wrong to waste a government student subsidy not doing any work, because of the personal neglect and the political ill consequences. Perhaps if either of these was absent and the other present, then it wouldn't be strictly *wrong* to waste a subsidy but rather merely objectionable. So we should think of this as a case where the justification for the wrongness is the *combination* of the distinct goods/bads. So there is one and only one reason, but that reason has an interesting internal structure.

Finally, there are multiple-reasons cases where each of the reasons appears to be doing distinct work. Say I should make myself quinoa salad as my dinner tonight. Perhaps an adequate justification for this is that it's particularly healthy, and I'll be hungry. That I'll be hungry establishes that I should eat *something*. That quinoa salad is healthy establishes that that I should eat *it*. It's not the case that I should eat quinoa salad even if I'm not hungry.

We can think of this as a case where there are multiple burdens of justification, and each reason answers to its own burden of justification. The

overall burden of justification is: why should I eat quinoa salad as my dinner tonight, rather than not eat anything as my dinner? There are two burdens of justification contained within this one: why should I eat dinner tonight rather than not; and why should I eat quinoa salad for dinner rather than something else? That I'll be hungry establishes that I should eat something for my dinner tonight. That quinoa salad is particularly healthy establishes that I should eat it for my dinner tonight rather than something else. The contrastive view of burdens of explanation, then, enables us to divide up the justificatory work in this way.

If this treatment of these various cases are correct, then we can simply construe reasons as being themselves justifications: justifications which answer particular burdens of explanation with their particular contrast classes. That was so in each of the above cases we encountered. This gives us a very simple equivalence.

One consequence of this equivalence is that normative reasons require a fact or facts which serve as the reason, and cannot rest on a false proposition. This is because a reason is a justification, a justification is an explanation, and an explanation requires a fact or facts which serve as an explanation. If the Colonel did not possess a gun, and knew it, then it cannot be part of the explanation of how he murdered the Marquis that he knew he could shoot him and avoid a precarious swordfight. In the same way, if quinoa salad isn't healthy, I can't justify my eating it from its being healthy. So we might think of this as an explanation of why reasons require facts: reasons require facts because reasons are explanations and explanations require facts. Sometimes reasons and explanations utilise things rather than facts, as when we say that the explanation of Britain's temperate climate is the Gulf Stream, or when one ought to refrain from bullying people because of the pain it causes them. A similar point applies here: the things which serve as the reason or explanation must exist. If there were no Gulf Stream, it couldn't possibly serve as an explanation of Britain's temperate climate. Similarly for normative reasons.

With this equivalence in hand, we can address the problem I raised through Ross in chapter 2. There, I argued that the paradigmatic problem of construing the role of reasons was that in some cases there are multiple reasons in play which genuinely conflict, in the sense that they recommend competing verdicts and the conflict is not internally resolvable, yet in at least some of those cases, the conflict between the reasons is resolvable nonetheless: there is a fact of the matter as to what the correct verdict is. The challenge, then, is to say in virtue of what the conflict is resolvable. It was by addressing a challenge of this kind that a weighing model appeared to gain traction.

Given that reasons are explanations, there are actually two problems here. One is how there can be reasons in favour of something that isn't true. For example, the possibility of taking it easy is a reason in favour of its being permissible to waste a government student subsidy doing no work, but still, because of the other things at stake, it's not actually permissible to do so. There can't be explanations of something that isn't true or that didn't happen: if there is no food, then no sense can be made of a putative explanation of why there is lots. There can be no such explanation. So one problem is how reasons can do this, given that reasons are explanations, and explanations can't do this. We must acknowledge this problem just so long as we acknowledge the above kind of case: cases where there are conflicting reasons one of which counts in favour of the truth a proposition that isn't in fact true. The second problem is the more traditional one: showing how that conflict could be resolvable, given that there can be conflicting reasons.

I will address the first problem first: if reasons are explanations, and there can be reasons in favour of false propositions, this entails that there are explanations of things that aren't true: and this appears to be an absurd consequence.

Instead, we should refine the equivalence of reasons and explanations. Instead of saying that reasons are explanations, we should think of reasons as items which would explain a normative truth were it a genuine truth. For example, the tedium of exercise is a reason in favour of the truth of the proposition that I am ill-advised to do it. What this means is: were it true that I am ill-advised

to exercise, one thing that at least partially explains this is the tedium of exercise. Putting this schematically: a reason in favour of P would explain P were P true (and does explain P if P is actually true).

An interesting objection to this is the following. Some think that normative truths are necessary, at least when shorn of all their empirical add-ons⁴⁵ (I don't mean to endorse this; rather, an interesting objection emerges if it is true). It is not hard to see the appeal of this view. It doesn't seem that it could be true, or that it could have been true, that torturing animals for fun is permissible. So it seems to be a necessary truth. The problem with necessary truths is that they result in impossible antecedents when counterfactualized. And from an impossible antecedent anything follows (though of course not everything follows from a necessary antecedent). So to say of a normative proposition P that if it were true, X would explain it, therefore implies that, necessarily, in the case were P is in fact necessarily false, if it were true X would explain it. Since the antecedent is impossible, it's trivially true that if P were true X would explain it, for any X. So anything at all could explain P were it true, and so anything at all counts as a reason for it, by my definition of a reason. So we have the consequence that absolutely everything is a reason for torturing animals for fun. What a disastrous result.

Something must have gone wrong with the metaphysics or the logic here. Still, I do not think we have to deny the necessity of norms in order to hold on to our counterfactual definition of a reason. There are conditionals with impossible antecedents that aren't trivially true. Whatever account is to be given of such conditionals can simply be exported to the normative case.

A nice example of a non-trivially true conditional with an impossible antecedent can be found in the mathematical/computer science problem P vs NP. This problem concerns whether for certain mathematical problems where putative solutions can be quickly verified, there is a quick algorithm for finding those solutions. My understanding is that the consensus is that there is no quick algorithm for finding those solutions, though a proof of this has not yet been discovered. So why is there a consensus? The idea is that if there

⁴⁵ E.g. p.31 of Thomson (1990)

were a quick algorithm for finding these solutions, it would have been discovered by now. No quick algorithm has been discovered by now. Therefore, there is none. The major premise in this syllogism is the conditional: if there were a quick algorithm, it would have been discovered by now. It is believed that the antecedent is necessarily false. But since the whole conditional is meant to be serious evidence for a mathematical truth, the whole conditional can't be trivially true. So this is an example of a non-trivially true conditional with an impossible antecedent. I expect that whatever account can be given of this is one that I can use with respect to reasons.

Now we can turn to the second, more central problem. The reasons for a normative proposition P are those that would explain P were P true, and the reasons against P are those that would explain P 's falsity, were P false. But reasons are supposed to tell us whether a normative proposition is true or not: normative propositions are supposed to be justified on the basis of the reasons. As it stands, then, my definition of reasons does not give us any tools with which to resolve deadlock.

At the end of the last chapter, I proved that it is sufficient for a non-terminal normative proposition to be fully justified (and hence true) that an adequate response can be given to all of its contrastive burdens of justification. The problem above is posed in terms that make clear that each of the competing propositions in question, P and $\neg P$, are non-terminal: if they are true, then their truth can be given a deeper explanation with reference to some other truths. If they are non-terminal, they are therefore subject to the proof. Since it cannot be that both are true, it cannot be that both are fully justified. Hence it cannot be that all the contrastive burdens of justification for both can be met.

This resolves the deadlock. Where there are reasons for competing normative propositions, then only one can be fully justified and hence true. That one is marked by its ability to meet all of its contrastive burdens of justification. This ability to meet all contrastive burdens of justification is what resolves

the deadlock. The conflict is resolved in favour of the proposition which can meet all of its burdens of justification.

Let's see how this goes in an example. Regular exercise increases one's life expectancy and improves one's immune system. These are reasons to do it, but more importantly reasons in favour of the truth of the proposition that one should do it, in some relatively undemanding sense of 'should'. Regular exercise, however, is tedious. That is a reason not to do it, but more importantly a reason in favour of the truth of the proposition that one shouldn't do it, again in an undemanding sense of 'should'.

It is in fact true that one should exercise. Why? What explains why one should exercise, rather than not, is that exercising increases one's life expectancy and improves one's immune system. If it were true that one shouldn't exercise, the tedium of exercise might explain this. But what explains why one should in fact exercise regularly rather than not is that it increases one's life expectancy and improves one's immune system. The tedium of exercise fails to explain why one shouldn't exercise rather than not. The former case explains all of its contrastive burdens, whereas the latter doesn't.

Clearly this is completely uninformative, and indeed it fails to answer the question of how the deadlock gets resolved. What we want to know is why an improved life expectancy and immune system, together, should establish that one should exercise regularly, rather than the tedium of exercise establishing the opposite. Both are genuine reasons, after all.

At last this gives us the real burden of justification, and it is important to notice how it adheres to the hallmarks of burdens of justification established in the last chapter. There is a common ground: there are genuine reasons on both sides of the argument. There are contrasting effects: one set of reasons establishes the conclusion it favours, while the set of reasons in favour of the opposite conclusion fails to establish it. The burden of justification itself has an explicitly contrastive form: why the one set of reasons should establish the conclusion it favours, while the other fails to.

So let's pursue the example down this road: we are to find a relevant difference between the tedium of exercise, on the one hand, and the prospect

of an improved life expectancy and immune system, on the other, which explains the difference in effect. Perhaps there are many relevant differences: let's focus on one for the sake of space. One can easily learn to cope with tedium, but a diminished life expectancy and immune system are near-irreparable losses (life expectancy in particular). So that's why the prospect of an improved life expectancy and immune system should establish that one should exercise regularly, while in contrast, the tedium of exercise should fail to establish that one shouldn't.

Reasons are explanations and explanations of normative truths work contrastively. So reasons are themselves contrastive. To say that the tedium of exercise is a reason that counts in favour of the truth of the proposition that I shouldn't do it, is to say that if it were true that I shouldn't exercise, then its tedium explains why I shouldn't rather than should. Furthermore the tedium of exercise would explain why I shouldn't exercise in contrast more generally to activities which aren't tedious, and where it's not the case where I shouldn't do them. Listening to music isn't (for the most part) tedious, for example, and it's not something I shouldn't do. Suppose we thought we shouldn't exercise, then a burden of justification might arise over why I shouldn't exercise but it is OK to listen to music. The tedium of exercise would then lift this burden.

It is worth outlining in succession the abstract differences between this view of the structure of justifications and the view of that structure that is contained in the idea of the weighing model. There are three principal differences. First, in that view the weightiest set of reasons established its conclusion. Since every reason has a positive weight, it follows that in cases where there is one reason in favour of P and no reason against P, then P. There is no such condition in my account: a reason is at least part of a justification for P, but where it is only part of a justification for P, and there are no other reasons, then there is no justification for P. Hence it is not a condition on being a reason that that reason establishes its conclusion unless it is defeated by other reasons. Perhaps there can be a version of a weight-based framework where this condition is modified, so this does not mark out a universal difference.

Second, if a reason has a certain weight or if it is weightier than another reason(s), this is something that has to be justified. But in contrast, on my account what has to be justified is only the complete relation of why some set of reasons should establish the conclusion it favours while a competing set of reasons fails to establish the conclusion it favours. So there is a shift in the burden of justification.

Thirdly, my account is not outwardly deductive. It does not deductively follow from the fact that a diminished life expectancy and immune system are near-irreparable losses, whereas doing something tedious isn't, that the former two establish that one should exercise regularly whereas the latter doesn't. Whereas, if these were associated with certain weights and were weightier than the latter, it would deductively follow that one should exercise regularly.

One thing this discussion should have brought out is that the proposition that X is a reason in favour of P is itself normative. As such it is open to the same kind of demand for justification as any other normative proposition. For example, we might imagine a worker who notices that whenever they are struggling with work, their boss is indifferent, while whenever a particular co-worker is struggling with work, their boss is eager to help. The following complaint to the boss makes sense: why should their co-worker's travails count as a reason for the boss to help them, while their own travails don't? This is a clear case of a burden of justification introduced by contrasting normative propositions, one which reasons are required to ameliorate.

Not only is the identification of reasons itself normatively contentious and open to the same kinds of justificatory demands as other normative propositions; burdens of justification also arise between situations where there are similar reasons at play and yet different normative verdicts are assigned. For example, suppose General X is deciding which enemy fortification to attack out of A and B. In favour of A is its coastal location which makes it a natural base for resupply and reinforcements. In favour of B is its symbolic importance in terms of the war objective; capturing it will raise morale and may lead to an earlier end to the war. Suppose General X

decides that he should attack B. Earlier X faced a similar decision as to where to attack where the candidates were two different enemy fortifications, C and D. As with A and B, C is coastal whereas D is symbolically important. Suppose X had decided to attack C in the earlier decision scenario. Now a burden of justification arises: why is B a worthier target than A, but C worthier than D, given that the reasons that bear on each situation are the same? And this burden can only be met by introducing an additional difference between the two scenarios which would explain the contrasting selection of targets. This difference would, of course, just disrupt the original analogy between the two scenarios.

Contrastive justifications intrinsically relate two separate scenarios and explain the difference between them: this is in their very form. In this way contrastive justifications count as an intrinsically systematic kind of ethical thought. We should think of this as a strength of the theory: it accounts for the natural tendency to think of ethics as relatively systematic. Even those who are skeptical about wholly or narrowly systematic ethical theory still acknowledge that ethics is to some extent inherently systematic. For instance, Bernard Williams writes:

‘A limited benevolent or altruistic sentiment may move almost anyone to think that he should act in a certain way on a given occasion, but that fact does not present him with the ethical...The ethical involves more, a whole network of considerations...’⁴⁶

And Williams acknowledges the intuitive force of the contrastive demand for further explanation of contrasting ethical judgements where the cases are supposedly similar:

‘[Someone is] inconsistent and irrational, if he counts intelligence and reliability as supposedly sufficient grounds for hiring a man and refuses to do so when considering a woman.’⁴⁷

A second advantage of my account is that it can account for the platitude mentioned in the last chapter that explanation enables understanding of what is explained. We can do this via relatively loose ideas of distinctiveness. Understanding why a truth holds is knowing what is distinctive of the truth,

⁴⁶ P.28 of Williams (2011), *Ethics and the Limits of Philosophy*

⁴⁷ P.128, *ibid.*

or a particular type to which it belongs, in virtue of which it should hold. For example, to understand why it should be true that breaking promises is wrong it is at least sufficient to understand the nature of promises and how this contributes to promises' being morally distinctive.

Because the way in which a given case contrasts with others is written into the very form of a contrastive reason-explanation, contrastive reason-explanations disclose the way in which their normative subject-matter is distinctive. For example, take the variety of contrastive questions that would arise in the case of deciding whether to vote for a particular candidate in a parliamentary election. Answering all those would involve knowing what all the relevant differences between the candidates were, the relevant differences between general elections and other types, and so on. Because giving a full contrastive explanation of why a particular normative truth holds contributes to knowing what is distinctive about its subject-matter, this makes clear how understanding the explanation yields understanding of the phenomenon.

Bibliography

- Anscombe, G.E. (1967). Who is Wronged? *The Oxford Review*, Number 5: Trinity
- Audi, R. (2004). *The Good in the Right*. Princeton University Press.
- Broome, J. (2013). *Rationality Through Reasoning*. Wiley-Blackwell.
- Broome, J. (2004). *Weighing Lives*. OUP
- Crisp, R. (2014). Moral Testimony Pessimism: A Defence. *Aristotelian Society Supplementary Volume* 88 (1):129-143
- Davidson, D. (1963). Actions, Reasons and Causes. *Journal of Philosophy* 60 (23):685-700
- Harman, G. (1986). *Change in View*. MIT Press
- Hills, A. (2009). Moral Testimony and moral epistemology. *Ethics* 120 (1):94-127
- Horty, J. (2012). *Reasons as Defaults*. OUP
- Johnson, Robert, "Kant's Moral Philosophy", *The Stanford Encyclopedia of Philosophy* (Summer 2014 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2014/entries/kant-moral/>
- Lipton, P. (2004). *Inference to the Best Explanation* 2nd edition. Routledge.
- Moretti, Luca and Piazza, Tommaso, "Transmission of Justification and Warrant", *The Stanford Encyclopedia of Philosophy* (Winter 2013 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2013/entries/transmission-justification-warrant/>.
- Nickel, P. (2001). Moral Testimony and its Authority. *Ethical Theory and Moral Practice* 4 (3):253-266
- Pears, D. (1998). *Motivated Irrationality*. Clarendon Press.
- Prichard, H.A. (1912). Does Moral Philosophy Rest on a Mistake? *Mind* 21 (81):21 – 37

- Rawls, J. (2005). *Political Liberalism* expanded edition. Columbia University Press
- Ross, W.D. (1930). *The Right and the Good*. Clarendon Press.
- Ruben, D. (1992). *Explaining Explanation*. Routledge.
- Ruben, D. (1988). A Puzzle About Posthumous Predication. *Philosophical Review* 97 (2):211-236.
- Scanlon, T.M. (2014). *Being Realistic About Reasons*. OUP.
- Schroeder, M. (2007). *Slaves of the Passions*. OUP.
- Setiya, K. (2007). *Reasons without Rationalism*. Princeton University Press.
- Skorupski, J. (2010). *The Domain of Reasons*. OUP
- Thomson, J. (1990) *The Realm of Rights*. Harvard University Press
- Van Frassen, B. (1980). *The Scientific Image*. OUP.
- Williams, B. (2011). *Ethics and the Limits of Philosophy* 3rd edition. Routledge Classics.