

Authority and Reason-Giving

David Enoch¹

1. The Problem of Authority

Arguably, you have authority over your 7-year-old son. This means, perhaps among other things, that you can, by your mere say-so, create duties for your son. You just have to tell him to go to his room, and suddenly he is under a duty to go to his room, a duty that just until a minute ago he did not have. Suddenly, by not going to his room – something that until just a minute ago was perfectly permissible for him to do – he will be acting wrongly. How did this magic happen? Similar magic seems to be going on whenever someone possesses legitimate authority. The state arguably possesses such authority over its citizens (or at least some states do, on some matters, some of the time). My dean arguably possesses such authority over me (regarding some professional matters). And perhaps also the passenger who seizes the initiative in an emergency, and starts to give instructions thereby coordinating the actions of the more panicked passengers (thereby perhaps saving lives): At least after some time has passed and people have been following her instructions for a while, she seems to be able to do similar normative magic with her words.

The problem of authority may have a metaphysical side to it. Duties and wrongness seem to be serious things, perhaps a part of the furniture of the universe all the way out there in Plato's heaven, or anyway – even if not as serious as all that – still pretty serious indeed. And so, with Raz (2006, 2012), we can ask "Is it that easy to manufacture duties out of thin air?". But I do not think that this metaphysical

¹ Earlier versions of this paper were presented at the philosophy departments in Duke, Stirling, Edinburgh, Washington University in St. Louis; at the epistemology group in the Van Leer Institute in Jerusalem; and in a conference at the university of Graz. I thank the participants – and especially Wilfried Hinsch, my commentator at Graz – for the valuable discussion. I also received helpful comments from Matthew Braddock, Antony Duff, Simon Hope, Ambrose Lee, Ofer Malcai, Arik Pelman, Stephen Perry, Joseph Raz, Ben Saunders, Walter Sinnott-Armstrong, Roy Sorensen, Gopal Sreenivasan, and Bas van der Vossen. Also, this research was supported by the Israel Science Foundation (grant no. 136/09).

puzzle about authority is all that puzzling, for reasons to be mentioned in section 3. A more pressing puzzle is the moral, or perhaps more generally normative one. Why should the mere say-so of one person make such a huge normative difference to another? This puzzle becomes more troubling still when we attend to the preemptive nature of authoritative directives (of which much more to follow): For when an authority issues a directive, the subject does not just get another reason for action he did not have until the directive was issued. Rather, the subject is in some sense expected not to act on his own discretion, perhaps not even to consider (some of) the other reasons for actions applying to him. He should, as it were, subject his will to that of the authority. But "how can it ever be that one has a duty to subject one's will and judgment to those of another?"² (Raz, 2006, 1012)

The topic has been receiving much scholarly attention, recently perhaps mostly in the context of critical evaluations of Joseph Raz's influential service conception of authority³. But I think that this literature has gotten us on the wrong trail, by (for the most part) failing to appreciate that the case of authority is a particular instance of a more general phenomenon, the one I call robust reason-giving. In this paper, then, I try to remedy this flaw. I think that by focusing attention on the more general phenomenon of which authority is a particular instance, and by relying on a theory of this general phenomenon, we can make progress on understanding authority as well as the recent debates about its nature and justification.

In the next section, I show that authority is plausibly considered a particular instance of robust reason-giving. Then, in section 3, I offer and briefly motivate a theory of robust reason-giving, a theory I defend in more detail elsewhere⁴. I then apply this theory to the case of authority in section 4. In section

² In what follows I will sometimes talk in terms of when an authority is justified, sometimes in terms of when an authority is legitimate. In some contexts it may be necessary to distinguish between the justification and legitimacy of authority (and in some contexts some people have). But I will use these locutions interchangeably: What I will mean by asking whether an authority is justified or legitimate is whether it is a genuine authority, one that can issue genuinely authoritative directives.

³ See Raz (2006; 2010); Darwall (2009; 2010); Sherman (2010); Herskovitz (2011); Viehoff (2011).

⁴ In my "Giving Practical Reasons" (2011).

5 I discuss one paradigmatic way of justifying authority, the one based on a duty-out-duty-in principle. In section 6 I discuss another possible way of justifying authority, one that does not depend on a pre-existing duty. The discussion in sections 5-6 places Raz's Normal Justification Thesis and the controversy over it in a new, theoretically productive, light. Section 6 is also where I discuss the phenomenon of preemption already mentioned above, and offer a preliminary account of it. Section 7 addresses a worry recently emphasized by Darwall and others in the context of criticizing Raz's service conception of authority – namely, that even if it highlights circumstances in which it will be rational for one person to act on the say-so of another, it misses the authority's *right* to rule, and so also the fact that the subject's duty to obey is *owed to the authority*, which is entitled to hold the subject accountable in case of a violation. I argue that talk about the right to rule should be debunked. Here I part company with what Raz says about this and related matters, though I think that my response is truer to Raz's underlying insights than Raz's own. Still, I agree that more needs to be done here, and so I offer a debunking explanation of the attractiveness of talk about a right to rule and the like. In section 8 I discuss in a preliminary way the relation between practical and theoretical authority, and in section 9 I briefly address the relation between authority and social practices, arguing that the availability of a relevant practice is not necessary for the existence of a legitimate authority, but that the account of authority in this paper nicely accommodates the contingent relevance of practices. In section 10, by way of conclusion, I offer a preliminary discussion of when authorities – understood as they are in this paper – are legitimate.

2. Authority as a Particular Instance of Robust Reason-Giving

We give each other reasons all the time, and it will be helpful to distinguish the kind of reason-giving relevant to authority from other kinds of reason-giving⁵.

⁵ Some parts in this section and the next one are taken from my "Giving Practical Reasons".

I have tentatively decided to tell a colleague exactly what I think of him, and it won't make for a charming scene. You urge me not to. I can then say something like "Give me one reason not to do it!". Suppose you reply by noting the bad effects such a scene will have on the intellectual atmosphere in our department. It seems like you succeeded in giving me a reason not to proceed with my ill-advised plan. But what you've done – the thing naturally described in terms of giving me a reason to shut up – is to *indicate* to me, or *show* me, a reason that was there all along, independently of your giving it to me. Perhaps, in my fury, I hadn't paid attention to it, and so your intervention can make a difference. But it didn't make a difference by way of creating a new reason. We can call such reason-giving *purely epistemic*, for the role of the giving here has nothing to do with the reason's existence, and everything to do with my knowing that it is there, appreciating it, and acting for it.

Now suppose your neighborhood grocer raised the price of milk. It is natural to say that she has thereby given you a reason to reduce your milk consumption. It is, after all, true that you didn't have this reason before her relevant action, that you do after it, and furthermore that you have this reason *because* of her raising the price. In a perfectly ordinary sense, then, she has created this reason, she has given you a reason to buy less milk. But there is nothing mysterious going on – no normative magic here. What the grocer did, it seems natural to say, is merely to manipulate the non-normative circumstances in such a way as to trigger a dormant reason that was there all along, independently of the grocer's actions. Arguably, you have a general reason (roughly) to save money. This reason doesn't depend on the grocer's raising of the price of milk. By raising the price of milk, the grocer triggered this general reason, thereby making it the case that you have a reason to reduce your milk consumption. Indeed, perhaps you even had all along the conditional reason to-buy-less-milk-if-the-price-goes-up. Again, this conditional reason doesn't depend for its existence on the grocer's actions. But the grocer can make the conditional reason into an unconditional one, simply by manipulating the relevant non-normative

circumstances. And this is what she did by raising the price of milk. I will call this kind of reason-giving *merely triggering reason-giving*.

But sometimes we manage to create reasons for action in what appears to be a more robust sense of reason-giving. For instance, if I promise you to meet you for lunch, I have just created a reason – indeed, a duty – for me to meet you for lunch. And the way in which I managed to create this reason seems importantly different from the one present either in cases of epistemic reason-giving or in cases of merely triggering reason-giving. Similarly, if I ask you to read my draft (a mediocre text about a topic you are not interested in, and so a text you do not have an independent reason to read), I can give you a reason to read my draft, and the giving here seems importantly different from that of purely epistemic reason-giving and from the merely triggering reason-giving discussed above. I give the name "robust reason-giving" to the kind of reason-giving arguably involved in cases of promises and requests.

It is robust reason-giving that is closely related to the phenomenon of normative magic I started with⁶. Presumably, there is nothing mysterious about the possibility of one person indicating to another the presence of a reason. And triggering a pre-existing conditional reason also does not seem mysterious at all – If the normative furniture all the way out there in Plato's heaven includes some conditional reasons with non-normative antecedents (and why wouldn't it?), then presumably there is nothing mysterious about the fact that we can trigger such conditional reasons by rendering their antecedents true. But as the literature on promises makes clear, there *is* something of a puzzle about the magical normative force of words when it comes to promises, and it seems to me clear – on phenomenological grounds, mostly – that something similar is going on in the case of requests as well.

⁶ Edmundson (2010) draws a distinction very similar to mine between merely triggering and robust reason-giving, using (following Estlund) the terminology of side-effect power and *really* exercising a moral power, respectively. I don't think this terminology is helpful – as some of the examples below show, triggering reason-giving may be intended, and so not a mere side-effect, and there's nothing less than fully real about the normative power to trigger reasons. Following Regan, Edmundson also distinguishes both these kinds from "epistemic reasons for action". This too is problematic terminology – what is epistemic here is not the *reason* (it is, after all, a reason *for action*), but rather the *giving* of the reason. But still, the distinction he has in mind seems to be the one in the text here.

Thus, once we have the tripartite classification of reason-giving above, it seems clear to me that the case of authority is supposed to be of this third kind, it is a kind of *robust* reason-giving. For when authorities issue directives what seems to be going on is the creation of a new reason, not merely the indication of a reason that was there all along (and this despite the fact that often the directives of at least good authorities will bear interesting relations to the reasons that are independent of the authority's directives; more on this shortly). And furthermore, the creation of reasons involved in authoritative directives is the kind present in requests and promises, not merely the one present in the raising-the-price-of-milk example.

So authority is an instance of the power to give reasons robustly. The sense of "power" here is roughly the Hohfeldian sense (Hohfeld 1917) of a normative power, the ability to bring about a normative change. Authority, it thus seems plausible to say, is an instance of the ability to bring about a change in the addressee's reasons for action. But crucially, the grocer also has this power – for when she raises the price of milk, she brings about a normative change in the reasons for actions that apply to you. Authority, though, is an instance of the power to give reasons *robustly*, that is in the special, seemingly magical way involved in requests and promises. But this still doesn't tell us in what ways the case of authority differs from that of requests (for instance). And here I want to introduce the following plausible hypothesis (which I will revisit several times in what follows): What's special about authority is that where it is present the authority can create not just any old reasons, but *duties*. So authority – real, legitimate authority – is the power to robustly give duties.

This is, as I said, a plausible hypothesis. But it needs further support. It will prove convenient to supply such support, though, only after we have an explicit account of robust reason-giving at hand. Let me turn to that, then, returning to authority only in section 4.

3. A Theory of Robust Reason-Giving

There are considerations in the metaphysics of reasons supporting the belief that robust reason-giving – if it ever occurs – can only be a particular instance of triggering reason-giving. This is so, roughly speaking, because whenever it's true that if you (say) ask that I ϕ I will then have a reason to ϕ , this conditional can only be plausibly explained by the antecedent existence of a reason to- ϕ -if-you-ask-me-to. And this means that the request triggered a conditional reason that was there all along. Similarly for other cases of robust reason-giving. (There is much more to say here – for a much more elaborate discussion of this metaphysical point, one that at the end of the day vindicates the point just made, see my "Giving Practical Reasons"⁷.)

Of course, it is at this point still a possibility that robust reason-giving is an illusion. But I want to note now that if it is not, then given the point in the previous paragraph, it should be clear why there is no serious metaphysical worry about robust reason-giving – in the case of authority, or elsewhere. For we already know that there is nothing metaphysically suspect about triggering reason-giving. And we also know that if there is such a thing as robust reason-giving, it is a particular instance of triggering reason-giving, and not, it seems, a *metaphysically* special particular instance either (though it is special in other ways). So there is no cause for metaphysical concern about robust reason-giving: At least, that is, if a phenomenon close enough to the pre-theoretical one of robust reason-giving can be accommodated as an instance of triggering reason-giving. Can this be done?

A promising first step in attempting to characterize robust reason-giving as an instance of triggering reason-giving is focusing on the intentions of the reason-giver. For clearly, one can give reasons in the epistemic or the mere triggering senses without even intending to thereby give a reason. But this is not the case for promises, requests, and commands. The giving of the reason is a part of the *point* of the relevant action, the reason-giver intends to give the reason-receiver a reason.

⁷ There I also discuss – following Schroeder (2005) – another way of explaining the truth of the conditional in the text, along the Constitutive Model. But this model is not relevant to the case of authority, except for the footnote below on divine authority.

But we shouldn't conclude too quickly that this is what is unique about robust reason-giving, for one can intentionally indicate the existence of a reason, or intentionally trigger a reason without giving a reason robustly. Suppose, for instance, that I put my foot on the road with the intention of thereby giving you (the driver) a reason to stop; still, in this case too what seems to be present is reason-giving of the merely triggering kind, not of the robust kind (as when I, an authorized police officer, tell you, the driver, to stop).

What seems to be missing in the case of the intentional merely-triggering reason-giving is not the intention to give a reason (for *this* intention is present in the pedestrian case), but rather the intention to give a reason *merely by the very forming of the intention to give a reason*⁸. Based on such considerations (and many more examples; again see "Giving Practical Reasons") we can offer the following account of robust reason-giving:

One person A attempts to robustly give another person B a reason to ϕ just in case (and because):

- (i) A intends to give B a reason to ϕ , and A communicates this intention to B;
- (ii) A intends B to recognize this intention;
- (iii) A intends B's given reason to ϕ to depend in an appropriate way on B's recognition of A's communicated intention to give B a reason to ϕ .⁹

⁸ This way of putting things makes the content-neutrality of such reason-giving more explicit. Such content-neutrality is often noted in the context of discussions of authority, but not in the context of requests. For a discussion that explicitly ties content-neutrality to intentions of the sort presented in the text, see Sciaraffa (2009).

⁹ There is an obvious structural similarity between this account and (one version of) Grice's account of sentence-meaning. It's not clear to me whether the similarity is merely superficial. I say a little more about this similarity in "Giving Practical Reasons".

For a similar understanding of requests, this time inspired by Searle rather than by Grice, see Cupit (1994, 450): "To request is to attempt to affect another's actions, by doing no more than presenting those wishes in a form which constitutes an attempt to affect action." Similarly, Raz (1975, 83) writes: "A person who makes a request intends his making the request to be a reason for the addressee to comply with it." And for explicitly Gricean discussions of closely related issues, see Hart (1982), and Sciaraffa (2009).

The third condition can be understood as a generalization of such natural thoughts as that when I ask you to ϕ , I intend that your reason for ϕ -ing be *that I asked you to*; that when I command that you ϕ , I intend that your reason for ϕ -ing be *that I said so*, etc.

Thus, when I ask you to read my draft, I intend to give you a reason to read my draft, and I communicate this intention to you by way of saying something like “Do you mind having a look?”¹⁰ (i); I intend you to recognize this intention (ii); and I intend this recognition of yours to play an appropriate role in your practical reasoning¹¹ (iii), as can be seen from the fact that my request misfires when you proceed to read the draft (as I asked) only for the reason that this will keep our department chair happy. In such a case you did as I asked, but not for the reason I intended you do it for. A similar analysis seems plausible for other cases of robust reason-giving.

There are two kinds of success conditions necessary for robust reason-giving. The first kind is non-normative: For A's attempt to robustly give B a reason to ϕ to succeed, B must recognize A's above specified intentions, and furthermore B must allow these intentions to play an appropriate role in his practical reasoning¹². Notice that this condition is not necessary for the attempt to succeed in amounting to a robust reason-giving, but rather for it to succeed in having the intended kind of effect in the world.

But this condition is not sufficient for the attempt at reason-giving to succeed, not even for it to succeed in amounting to a robust reason-giving. To see this, think of the example of the dictator's

¹⁰ The communication is important here. It is not sufficient for robust reason-giving that I intend that *you know* about my relevant intentions; it is important (as can be seen from (iii)) that what plays a role in your practical reasoning is that I actually communicated to you the intention that you take this very communicated intention as a reason. (I thank Joseph Raz for emphasizing this point to me.)

¹¹ As I am about to explain, there are two success conditions relevant for the attempt to robustly give a reason, one normative, one non-normative. Because both are required, I can use wording that is not explicit about the distinction between normative reasons (as in (iii)), and motivating reasons (as in the text here).

¹² Notice that the “communicates” in condition (i) is not understood as a success term, requiring uptake of some kind. It requires merely the attempt to communicate the relevant intention. (I thank Hanoch Sheinman for a related point).

child¹³: The son of a brutal dictator "orders" you to perform some action. He is, of course, not authorized to issue such an order, not even according to the rules his dictator father accepts. But if you don't do as the child says, he will become cranky; and if the child becomes cranky, his father will be cranky; and if the dictator becomes too cranky, he will brutalize some innocent people. In this case, it seems you now have a reason (indeed, an obligation or a duty; but for my purposes a reason will do) to do whatever it is the dictator's son ordered you to do. The dictator's son has succeeded in giving you a reason for action. And indeed, this is exactly what he intended to do. This, of course, is not a case of robust reason-giving. It is merely a case of a triggering reason-giving: You have a standing reason to prevent horrible disasters from befalling innocent people, and the dictator's son has successfully manipulated the non-normative circumstances so that this general reason will imply a more specific one to do as he says. But if he doesn't have the intentions characterized in clauses (i)-(iii) above, the dictator's child doesn't even qualify as *attempting* to robustly give you a reason for action. Now assume that the dictator's son does have all of the intentions specified in (i)-(iii) above, so that he genuinely does attempt to robustly give a reason; and assume further, that the person he addresses treats him as genuinely authorized or in some other way able to robustly give reasons in this way, so that she lets the child's reason-giving intentions play the appropriate role in her practical reasoning. Still, it's clear that the child has not managed to robustly give a reason for action. What is missing in this case is the *normative* success-condition, namely, the attempt must make it the case that a normative reason to ϕ really does emerge (in the appropriate way). And we already know that whether this procedure will result in there being a reason to ϕ here will depend on there being an independent reason that is triggered by this procedure – roughly, a reason (for B) to do as A intends that B have a reason to do. In the dictator's child example, there is no general reason to do as the child "commands", and so his reason-giving intentions do not trigger such a general reason, and so he does not succeed in robustly giving a reason, whether or not his audience believes

¹³ I use it in "Giving Practical Reasons", where I take it from Estlund (2008, 118), who in turn credits John Deigh for the example.

that he does. But when I ask you to read my paper, presumably there is this general reason (to do as I ask, within limits, in a certain context, etc.), one that I presumably succeed in triggering by making the request. It is in this way, then, that the suggested account of robust reason-giving is a particular instance (but an importantly unique one) of triggering reason-giving. And we can now say that for robust reason-giving to occur, there must be, independently of the attempt at robust reason-giving, a reason triggerable by such an attempt.

I am not sure what more to say about the "appropriate way" qualification in (iii). It is meant to rule out deviant causal (and perhaps other) chains. It would have been nice to have an explicit account of how exactly to do this. But I will have to settle for noting that usually we know a deviant causal chain when we see one, and for claiming partners in guilt – for almost everyone needs an account of deviant causal chains. This qualification in (iii) thus doesn't make (iii) (or the account of which it is a part) empty, nor does it raise any new problems that are peculiar to my account of robust reason-giving. I should also note something it does *not* take for the role played by the given reason in the receiver's practical reasoning to be appropriate. It is not required that the role be, as it were, ultimate¹⁴. In other words, it is perfectly consistent with robust reason-giving thus understood that there be a further, fuller, perhaps more basic story of why it is that B does and should take A's relevant intentions as reason-giving. Perhaps, for instance, B is a simple utilitarian, and let's further assume that simple utilitarianism is indeed the true fundamental story about all reasons for action. If so, she will take A's request as a reason to ϕ if and only if, and because, doing so will maximize utility. But this does not mean that she doesn't take – in those cases – A's request to be a (non-ultimate) reason. The crucial question is whether the ultimate (or perhaps just more basic) story here is one that goes through the reason-giver's special intentions identified above (and the receiver's recognition thereof), as in the case of the utilitarian request-receiver, in which case we may have a case of robust reason-giving; or whether the more basic

¹⁴ I thank Cian Dorr for pressing me on a related point.

story here works directly, leaving no role for the specific intentions that make reason-giving robust (as is the case in the dictator's child example). Cases of this latter type are not, on the account I'm suggesting here, cases of robust reason-giving. And this seems to me the independently plausible result here¹⁵.

So much, then, for robust reason-giving in general. It is time to get back to the particular instance of authority.

4. Authority, Again

In the opening paragraph, you told your son to go to his room, and you did this, we are assuming, with authority. I think we can see that the conditions of robust reason-giving from the previous section are met: (i) you intended to give your son a reason to go to his room, and you communicated this intention to him; (ii) you intended him to recognize this intention of yours; (iii) you intended this recognition to play an appropriate role in his practical reasoning, that is (roughly speaking) you intended him to go to his room for the reason that you told him so to do. And the normative success condition was also met – for presumably your 7-year-old has a reason to do as you say (on some matters, within limits, on some occasions, etc.), indeed perhaps he's had all along the reason to-go-to-his-room-if-you-tell-him-to, a reason that you successfully triggered precisely by telling him to go to his room. So the above analysis of robust reason-giving seems to apply to this case of authoritative reason-giving, and, it seems reasonable to hypothesize, to others as well.

¹⁵ But – as a conversation with Antony Duff helped me see – things are more complicated here. What should we say of a sufficiently *systematic* variation of the dictator's child case, where there is now a practice of "obeying" the child, etc.? First, it's possible that if the story is sufficiently beefed up, at some point it will become a story of genuine authority. Second, and relatedly, whether there is here legitimate authority or not may be a matter of degree. Third, it's possible that the only way for a consequentialist to avoid implausible claims here would be to insist on some *general* duty to obey; of course, we can describe a version of the dictator's child case that is general in such a way. About such a case, a consequentialist may have to settle for the two first points above. (This will get her closer to the practice-dependant theory of authority I criticize below, in section 9, but not quite all the way there.) Fourth, *some* kind of generality may be involved even regardless of consequentialism, as it is arguably a part of the Razian preemption story, to which I get later on.

But all this would have been true also if you *asked* your son to go to his room, for requests too are a case of robust reason-giving. What makes it the case that you *told* him, rather than *asked* him to go to his room? Surely, not just the wording you chose to use, as we can use wording that looks superficially like a request to issue a command or command-sounding wording to make a request. What does the work here, I think, is the plausible hypothesis from section 2, namely, the thought that the distinctive feature of authorities is that the reasons they can give robustly are reasons of a special kind, *duties*. The point is not, of course, that the reason-giver merely *manages* to give the reason-receiver a duty. This can happen in cases where no authority is present. For instance, if I request that you read my draft, and for some reason you have an independent duty to accommodate my request (perhaps because I've done so with your requests often enough in the past, or because you promised our department chair you would, or some such), then I managed to give you a duty, but without possessing authority here; and the pedestrian placing her foot on the road manages to give the driver a duty – not just a reason – to stop. What is needed for authority is that the intentions specified in the three conditions (i)-(iii) for robust reason-giving apply to duties¹⁶. So we get:

One person A attempts to command another person B to ϕ (or, in other words, claims authority¹⁷ to tell B to ϕ) just in case (and because):

- (i) A intends to give B a duty to ϕ , and A communicates this intention to B;
- (ii) A intends B to recognize this intention;
- (iii) A intends B's given duty to ϕ to depend in an appropriate way on B's recognition of A's communicated intention to give B a duty to ϕ .

¹⁶ Below, after discussing the role of protected and quasi-protected reasons, I will have to qualify this statement, and accordingly also the three conditions that follow in the text.

At one point, what Raz (1986, 54) says about *non*-normal justification of authority makes it sound as if he would classify at least some of the cases I would classify as the mere triggering of a duty as cases of genuine authority that is non-normally justified. I don't think this difference is important.

¹⁷ In one sense of this problematic phrase. There are others, I'm sure.

And of course, a normative success condition also applies. Only against certain normative backgrounds will A actually succeed in creating a duty for B to ϕ by forming and communicating these intentions. Let me not say more now about the nature of the normative success condition, because this will be our main order of business in the next two sections.

In the more general case of robustly giving *reasons* it is not a necessary condition that the reason-giver intend the reason-receiver actually to ϕ . This is so, because (as I explain in more detail in "Giving Practical Reasons") I can, for instance, request that you read my draft, intending to thereby give you a reason to read it, without also intending this reason to be conclusive or to defeat all others. If I don't know what other reasons apply to you, for instance, I can intend to give you a reason (by requesting), all the time realizing that if other countervailing reasons apply and outweigh the one I just gave you, you won't read my draft. And I need not intend that this not be so. Of course, if I *know* that you have stronger reason not to read my draft, there is something non-standard about my request, perhaps to the point of us refusing to classify it as a genuine request after all. But I may make a genuine, sincere request with the intention to make a difference to the reasons applying to you (and so also to their balance) without necessarily intending that the weight of reasons will ultimately support you doing as I request. Now, I am not sure, but here we may have a difference when it comes to robustly giving *duties*. For it does seem at the very least odd that someone would issue a command that the addressee ϕ without also intending that she ϕ . The way in which giving a reason to ϕ is consistent with acknowledging the possibility of stronger countervailing reasons does not seem to apply to the case of duties – robustly giving a duty does not seem consistent with acknowledging that the addressee may after all have stronger (normative) reasons not to comply. But I do not think that this entails a strong disanalogy between the duty case and the more general reason-case, or in some other way threatens the analysis above, for the following two reasons. First, contrary to first appearances, there may be analogous cases in the duty case as well. Perhaps, for instance, a state official can genuinely and

sincerely issue a directive all the time acknowledging that there may be circumstances (perhaps, for instance, circumstances involving a conflict of duties) in which her directives ought not – all-things-considered – be obeyed. Of course, if the official *knows* that the addressee has stronger reason not to comply, there is something non-standard about the directive, perhaps to the point of us refusing to classify it as a genuine authoritative directive after all¹⁸. But as we saw, this is true in the reasons case as well. Second, if there is a difference here between the reason- and the duty-case, I think it is due to the ways in which duties are special among other reasons. Roughly, because duties are typically especially strong, and perhaps also protected¹⁹, reasons, so naturally, an intention to give a reason of this kind will not be as easily reconcilable with an absence of an intention that the addressee actually act accordingly.

So far I've fleshed out some of the details of the plausible hypothesis that authority is the power to robustly give duties, but I haven't *argued* for it. Let me change this now. We have two directions. Assume first that someone A has the Hohfeldian power to give B duties, in the special way that is involved in robust reason-giving, namely (if the analysis presented in section 3 is right) by forming the intentions in (i)-(iii). It seems clear to me that in such a case, A intuitively has authority over B. For A can – by his mere say-so, or more precisely merely by forming and communicating the relevant intentions – make it the case that B has the relevant duty. The account suggested thus offers a simple vindication of the thought that an authority can create duties merely by uttering certain words (backed by certain intentions). Now suppose (for the other direction) that A is in a position of authority over B with regard to matters such as ϕ -ing. This must mean, at least, that A has it within her power to make it the case that B is under a duty to ϕ . The Hohfeldian power to create a duty is, in other words, a necessary

¹⁸ There may be interesting borderline cases here. If someone utters what sounds like a command not caring whether the subject acts accordingly, just in order to protect himself from being criticized for having failed to issue that command, is she genuinely commanding? (I thank Joseph Raz for this example; he answers in the positive.) If a police officer issues what sounds like an authoritative directive at a citizen which – as the officer knows – will not comply, is the officer genuinely issuing a command? (I thank Roy Sorensen for this example; he answers in the positive – certainly, the citizen won't be entitled to defend his non-compliance by noting that the officer wasn't genuinely issuing a command). Even if these cases are cases of genuine commands, they are in a sense non-standard, and perhaps also parasitic on the more standard cases.

¹⁹ There will be much more on protected reasons below.

condition for authority. But it is not sufficient, as duties – much like reasons – can be merely triggered. To repeat, when the pedestrian places her foot on the road, you – the driver – suddenly have a *duty*, not just a reason, to stop. But this is no manifestation of *authority* on the pedestrian's part, though it is a manifestation of one of her Hohfeldian powers. Here too, then, we need the distinction between merely triggering reason-giving and robust reason-giving. If A has authority over B with regard to matters such as ϕ -ing, then it's not just that A can give B a duty to ϕ ; A can give this duty robustly. And this is precisely what this direction of the plausible hypothesis asserts.

We have, then, a partial vindication of the initially plausible hypothesis: authority consists in the power to robustly create duties, with robustness fleshed out in terms of the intentions above²⁰. And this hypothesis gains further support from the fact that it nicely fits the examples of authority we started with, or so at least it seems to me. As a parent you do have the power to create – by forming the intentions mentioned above – a duty for your 7-year-old to go to his room. Most other people do not have such authority – even if they form the intentions (i)-(iii) the normative success condition is arguably not met, and so they do not have it within their normative powers to create by forming and communicating the relevant intentions a duty for your son to go to his room. Those of us who believe in the authority of the state (some states, some of the time, on some matters) accept that the state can create duties in its subjects, and that it can create them in the way that is present in promises and requests, not in mere-triggering cases²¹. And those of us who believe that the passenger who takes the initiative in an emergency possesses authority seem to believe something similar – indeed, those of us who *reject* this claim seem to reject it *because* they think that the passenger, while clearly being able to

²⁰ This has (as Gopal Sreenivasan noted) the somewhat awkward result that in cases of promises, we exercise authority over our (future) selves. I am not sure whether I am willing to accept this result, or whether I want to revise the understanding of authority so that it consists (roughly speaking) in the Hohfeldian power to robustly give duties *to another*. It doesn't seem to me that anything of substance depends on this choice.

²¹ For this to be so, it would have to be the case that the state can form intentions of the kind specified in the suggested analysis of robust reason-giving. I think that this is as it should be, and that the fact that the suggested account does not allow for robust reason-giving to be done by intention-less creatures is an advantage of it, as I am about to argue in the text.

create duties (if she shouts "everyone stay on the left!", and if people are likely to do as she says, and if the best way to prevent loss of lives is if everyone stays on the same side, then you have a duty to stay on the left), still creates duties by merely triggering them, not by robustly giving them.

Let me mention another explanatory advantage of this way of viewing authority. Pre-theoretically, it seems that only persons or person-like organs can play the role of an authority²². But not on all accounts of authority is it clear why this should be so. On an epistemic account, according to which authorities merely assist the subject in knowing about the reasons for action that she already has, it seems that anything that can serve such an epistemic function can be an authority. But this seems false – a barometer can indicate to me the presence of an overriding reason for me to take an umbrella, but it doesn't have authority over me. More interestingly, on Raz's service conception of authority it is also not clear why only persons (or person-like things) should be able to possess authority. After all, other things can also help one better conform to reasons, and this – according to the Normal Justification Thesis, the centerpiece of the service conception of authority – usually suffices for legitimate authority²³. Now, Raz can add conditions to make sure that nothing too different from a person will have on his theory legitimate authority. But if the only rationale Raz can offer for such

²² Don't other things sometimes possess authority? Don't traffic lights? Doesn't *the law*? At least, isn't it the case that it could? Much more needs to be said here, of course. But let me just state (without argument) the following: For some of these cases, it may be that we are happy to attribute to them authority only in a very loose sense (perhaps that's the case with traffic lights). And when something that is not a person or even person-like possesses genuine authority – if that is ever the case – this is so *in virtue of* some person or person-like thing (like perhaps the legislature) possessing authority. I thank Simon Hope and John Doris for pressing on me the question briefly discussed in this footnote.

²³ Raz often (e.g. 2006, 1017) uses the example of an alarm clock given you protected reasons. And Darwall (2010, 270 and on) uses versions of this example as central parts of his recent criticism of the service conception of authority. I will return to some of his relevant points there later on in the text. For now, though, let me note the following: On my suggested way of viewing authority as a particular instance of robust reason-giving, the problem with the alarm clock is not that it fails to give reasons for action (it succeeds in doing that), and not even that it fails to give preemptive or protected reasons for action (it may succeed in doing that, though see the discussion of preemption in section 6 below). The problem, rather, is that whatever reasons it does give it fails to give robustly, it merely triggers them. And this result seems to me exactly the right thing to say about this example (and we can achieve this result without going second-personal in anything like the way Darwall thinks is inevitable). Perhaps Darwall is misled here by his failure to take explicitly into account the distinction between robust and merely triggering reason-giving.

further conditions is to save his theory from counterexample, their introduction is of course objectionably ad hoc. However, once we see that the case of authority is a particular instance of robust reason-giving, and once we have an account of robust reason-giving in terms of the rather complex intentions involved, it just naturally falls out of the theory – without the need for any ad hoc further conditions – that only things that can form and communicate such complex intentions can have authority²⁴.

Before concluding this section, let me mention another theoretical payoff that viewing authority as a particular instance of robust reason-giving can secure. When it comes to authority, some people – libertarians, we can perhaps call them – believe that the consent of the subject is a necessary condition for authority. We cannot be subject to others' powers of this kind, the thought seems to be, unless by our own consent we have made ourselves so subject. But now that we know that authority is a particular instance of a wider phenomenon, there is something perplexing about this libertarian thought. No one thinks, for instance, that we can't be subject to others' power to robustly give us reasons unless by our consent. No one denies, as far as I know, and no one should, that a *request* from someone can give us reasons for action (robustly) regardless of our consent. So in order to render the libertarian thought at all plausible, it is just not enough to talk of man's liberty, or to reject involuntary subjection to others' will or power. In order to render the libertarian thought at all plausible, what we need is a story that explains why *duties* are special; why, in other words, even though there's nothing problematic about us being involuntarily subject to others' *reason-giving* power, we cannot be involuntarily subject to their *duty-giving* power. Perhaps such a story can be told, but until it is, this

²⁴ Another possibility is that Raz does not after all need a further condition restricting authority to just persons and person-like things, because he has always understood authority as a particular instance of (what I call) robust reason-giving, with this phenomenon understood roughly as I understand it here. And in fact, Joseph Raz patiently explained to me that this indeed is the case. I am not sure to what extent his texts easily lend themselves to such a reading, and I am quite confident they haven't always been read in this way. But needless to say, I am happy (and reassured) to think of what I am saying in the text as an elaboration of Raz's original ideas rather than as a criticism thereof.

influential libertarian thought just loses any plausibility whatsoever²⁵. (I return briefly to a discussion of consent in the final section.)

I do not want to create a false impression about what's been achieved so far. Even if everything so far said is true, all we have at this stage is a better understanding of what authority consists in. We still haven't even embarked on the normative task of seeing when an authority is justified or legitimate. But it seems plausible to suppose that now that we are better equipped conceptually, we should be better placed to start on that normative task. To this we now turn.

5. Duty-Out-Duty-In?

For an attempt at an authoritative command to succeed, the normative condition must be met. That is, it must be the case that antecedently to the purportedly authoritative directive and independently of it, things are such that if the authority issues a certain directive, then the relevant subject is under the relevant duty²⁶. Under what conditions will this be the case? When, in other words, is the claim to authority normatively successful?

In the more general case of robust reason-giving, my answer to the analogous questions was simple: An attempt at a robust reason-giving can only succeed if there is an antecedent reason that makes it successful. A request, for instance, could only succeed in robustly giving a reason for action

²⁵ The line of thought in the text applies also, I think, to views that may not be exactly libertarian in the sense explained in the text, but that are still close enough in the relevant ways. See, for instance, Edmundson (2010, 187 and 190), and Smith (forthcoming).

²⁶ In his "Political Authority and Political Obligation" (manuscript), Stephen Perry insists that a part of what it is for one agent to have authority over another is for it to be good – or to serve some value – that this is so (see especially his discussion of "the value of intentionality"). And Joseph Raz stressed to me the centrality of some similar condition to his thought of authority. I agree that there are *some* important relations between there being a legitimate authority and it being good that there is such an authority. Indeed, I believe there is such a connection in the case of normative powers in general (of which authority is a particular instance). See my "Being Responsible, Taking Responsibility, and Penumbra Agency" (forthcoming). But unlike Perry, I don't think that this is a part of what authority *is*; rather, it's a part of what in general justifies authorities. This leaves room – on my understanding, and I think also on Raz's, but not on Perry's – for specific cases in which the existence of a normative power and the desirability of its existence come apart (in both directions).

when there is a reason – independently of the request – for the addressee to do as the requestor requests. And so a *very* natural thought suggests itself in the case of authority as well. Perhaps what is needed for a successful issuing of an authoritative directive is a pre-existing *duty*, a duty that is independent of the authoritative directive and that requires, perhaps roughly, to do as the authority requires that one do.

I am going to raise some worries about this suggestion in a minute. But for now it's important to appreciate its strength. For clearly, if for some reason B is already under a duty to do as A says, then A possesses the Hohfeldian power to create duties – indeed, robustly – for B. And on the plausible hypothesis that authority consists in the power to robustly give duties, it follows that A is an authority over B. So it seems like we're in a position to conclude that whenever there's a pre-existing duty (roughly speaking) to obey, there is authority.

You may think that this gets things backwards. Aren't we interested in a theory of authority partly because we want to know when we should obey? Wasn't the (perhaps implicit) hope of the whole project to come to the discussion of duties to obey already equipped with a theory of justified authority? Wasn't the hope to show that at least sometimes we should obey *because* the authority is justified or legitimate, not that the authority is justified or legitimate because we already have a duty to obey, a duty that we presumably do not need the concept of authority at all to understand?

Several points should be made in response to this worry. First, and rather boringly, we sometimes use such "because"s as a way of explaining specific cases utilizing generalizations. It is presumably in this sense that we can say that the water boiled at 100 degrees centigrade because it was at sea-level, and this is the boiling temperature for water at sea-level. In this sense, then, there's nothing problematic in saying that your 7-year-old should go to his room because you told him to, and you have legitimate parental authority over him. This somewhat superficial explanation is consistent with your parental authority being ultimately grounded in his duty to obey rather than the other way around (just

like the superficial explanation above is consistent with the ultimate metaphysical relation going in the other direction, so that the boiling temperature for water is 100 degrees centigrade because this is when water at sea-level boils, and not the other way around). And of course, on some occasions, such superficial explanations are perfectly adequate.

Second, and much more importantly: So far all we have (pretty much) is the thought that a pre-existing duty to obey *suffices* for authority. Hardly anything has been thus far said to support the claim that such duty to obey is *necessary* for authority. Perhaps, in other words, though sometimes the normative success condition for authority is satisfied in virtue of an independent duty to obey, at other times it is satisfied in some other way. If so, then in those cases there may still be a deep sense in which legitimate authority is prior to the duty to obey rather than the other way around. I return to the possibility of other ways of generating a duty to obey later on.

But third, let me concede at least the following point: If the only way to satisfy the normative success condition for authority is for there to pre-exist a duty to obey then the theory of authority does relatively little explanatory and justificatory work. On reflection, I do not think that this is surprising²⁷. If it takes a duty to generate a duty, then authorities merely, as it were, channel the force of pre-existing duties to specific directives. This is not nothing, exactly. But it's true that on this understanding there's something less exciting about authority than you may have thought. In particular, an account of

²⁷ The point in the text is very closely related to a recent exchange between Raz and Scott Hershovitz. Raz (2006, 1030) has recently argued that cases in which there is an antecedent duty to obey (or something close enough to it), far from serving as counterexamples to his Normal Justification Thesis, are actually cases where the NJT is guaranteed to be satisfied, as in those cases one is virtually guaranteed to do better in conforming to reasons by obeying the authority than by attempting to act on the first-order reasons directly, as one's presumably most weighty reason relevant now is the duty to obey itself, and one best conforms to it by, well, obeying the authority. Hershovitz (2011) responds by noting that this way of understanding the NJT renders it void of any explanatory force. Now, I think that Hershovitz has a point here, but also that he overstates his case. He is right in that *for those cases* in which the NJT applies because of a prior duty to obey, the NJT itself (and the service conception with it) cannot explain the general duty to obey. But Hershovitz overstates his case, because not all cases need be of this kind, and because if the NJT does considerable explanatory work for other cases, the fact that it doesn't do this explanatory work for the cases of the kind discussed in this footnote need not count against it: In such a case, for anything thus far said the NJT may still be the best unified account of legitimate authority, doing interesting explanatory work with regard to some of the relevant cases, and settling for not being refuted by others.

authority does not then seem like a very good focal point for doing political philosophy – what is needed, rather, is an account of duties (including duties to obey); an account of legitimate authority would then fall out of it.

On this account, then, it takes a duty to generate a duty. So far I've been talking about a duty to obey, which – when the authority issues a directive, perhaps telling the subject to ϕ – yields the subject's duty to ϕ (on the say-so of the authority). But duties can serve as input also in another way. Think again of the helpful passenger who seizes the initiative in an emergency and manages (so we can assume) to coordinate the actions of all those present. If enough is at stake (if, say, non-coordinated action will result in a loss of lives), all present are under a duty to do what they can to coordinate their actions. This duty does not depend, of course, on the directives issued by the volunteer. But once she has spoken up ("Everyone, to the left!"), it becomes clear that the best, perhaps only, way to coordinate is around her instructions. So the duty that you have to act in a coordinated way now entails the duty to act according to the volunteer's instructions. Here too, then, we have a duty that backs up the relevant directives. But I do not think that here we have legitimate authority quite yet. True, the volunteer has the Hohfeldian power to create duties in this way (because of the background duty everyone is under to coordinate); but the way in which she can create these duties seems to me to be the merely-triggering way – it is very similar, for instance, to the dictator's child example. And the account of robust reason-giving from section 3 supports this point. For you do not have to know anything about the volunteer's intentions – in particular, whether she has the intentions distinctive of robust reason-giving – in order to know that you are under a duty to follow her instructions. Even if she is busy reciting lines from her favorite play, so long as her "instructions" are salient²⁸ in the relevant way, the best way to coordinate may be on her instructions, and then the duty to coordinate will translate into a duty to follow her "instructions". In this case, then, the pre-existing duty cannot satisfy the normative success condition for

²⁸ Lewis (1969).

authority, because the case is not a case of *robust* reason-giving at all. But time, it seems to me, can make a difference here. If the practice of coordinating on the volunteer's instructions has been going on for a while, we may now have a reason – perhaps even a duty – to obey her. Indeed, now her intentions *may* make a difference: Perhaps, for instance, as time goes on and the coordination becomes more systematic, we will want to distinguish cases of the volunteer's reciting lines from her favorite play and cases of her genuinely giving all of us instructions, and one natural way of doing that will be in terms of her reason-giving intentions. It seems to me that we can fill in the details in a plausible way so that regularity and systematicity will eventually suffice for making the difference between the volunteer's instructions giving reasons in the merely triggering sort of way and their constituting legitimately authoritative directives. In those cases too, what ultimately explains her status as authority is the need to solve coordination problems. But now it will do so via a general duty to obey her; and this will then suffice for authority. This result – naturally implied by the way I suggest we understand talk of authority and its normative success condition – seems to me to nicely accord with our reflective intuitions about the matter.

So much, then, for ways in which a prior duty to obey may ground legitimate authority. But it is now time to return to the question whether this is the *only* way in which a legitimate authority can be grounded. Can it be, in other words, that there is no prior duty to obey, that the authority is justified in some other way, and that then a duty to obey emerges (because, after all, it is the directive of a legitimate authority)? The question is not, of course, about temporal priority, but about (explanatory, and so in our context, where the explanandum is normative) justificatory priority.

You may be tempted by a duty-out-duty-in principle, a kind of conservation law for duties, according to which we cannot get a duty as the output of anything unless a duty was already there as a part of the relevant input. If you endorse a duty-out-duty-in principle, you will be inclined to think that authority must be grounded in a prior duty (most likely, to obey). For after all, an authority often has as

its output duties, and so must – given a duty-out-duty-in principle – have a duty as its input, and the only plausible duty in the neighborhood seems to be the duty to obey²⁹. But why accept a duty-out-duty-in principle?

It's important to distinguish between different conservation laws. I, for instance, accept a normativity-out-normativity-in principle, as I believe that normativity cannot be reduced to anything not itself clearly normative³⁰. And for somewhat related reasons, I defend a reason-out-reason-in principle for robust reason-giving in general³¹. But a conservation law for *normativity* is one thing, and a conservation law for *duties* is quite another. Perhaps even if the normative/non-normative divide is unbridgeable, this does not mean that the duty/non-duty divide is as unbridgeable. Perhaps, for instance, duties can be given a satisfactory account in other terms, themselves still normative. If this can be done, this will refute duty-out-duty-in, and will allow us then to ask in particular whether a duty to obey can be manufactured from other, non-duty though still normative materials.

6. Preemption

It is here, I think, that we should revisit both the preemptive phenomenology of authority, and Joseph Raz's service conception thereof.

Start with the former. When a legitimate authority issues a directive, it seems like what is (normatively) expected from the subject is not just to take that directive into account in his deliberation

²⁹ Let me emphasize here that the duty to obey need not be a duty to obey *a person*. The duty may be a duty to faithfully follow a procedure, say the democratic procedure. If you have such duty, and the procedure yields a directive, something like authority may be going on, and it is still grounded in the prior duty. I think that this observation suffices as a defense against the proceduralist objection to Raz's theory of authority (see Hershovitz (2011), and the references there; for a related discussion see also Viehoff (2011): Viehoff ends up saying something very similar to what I'm saying here, namely, that sometimes an authority may help a subject conform to reasons in virtue of the reasons the subject has to follow a certain procedure. But on the way he also seems to say that there is something in the NJT that is not easily reconcilable with the prior-duty disjunct of justifying authority (2011, 252). If he does say that, I differ, of course.).

³⁰ For an elaborate defense of Robust Realism about normativity – that entails such irreducibility – see my *Taking Morality Seriously: A Defense of Robust Realism* (2011).

³¹ Again see "Giving Practical Reasons".

as another factor to be weighed in with the others. Your son, when you tell him to go to his room, should not think to himself "Well, my favorite show is on, and the television is in the living room, so this counts against my going to my room; my favorite toy is in my room, so this counts for my going to my room; also, Dad wants me to go to my room, so this too counts for my going to my room; all in all, then, I should go to my room. Off I go.". Rather, he should think to himself something like "Dad told me to go to my room, and I should do as he says. Off I go." There are limits, of course, to the kind of preemption involved in even legitimate authority, but some preemption seems to be of the essence of authority – the subject is supposed to replace his or her discretion, to an extent at least, with the authority's directive. An account of authority that misses this aspect will seem to have missed an important part of the phenomenon to be explained.

Raz attempts to accommodate this preemption using the device of exclusionary and protected reasons. An exclusionary reason is a reason not to act for some reason or range of reasons. Suppose, for instance, that you're considering giving money to famine relief, and there are two reasons for doing so the strength of which you are currently feeling: that giving money will alleviate the suffering of those dying of hunger in third-world countries, and that giving money will make you look good to the colleague in whose attention you are interested. Both of these, it seems, are good (though perhaps not equally good) reasons for giving money to famine relief – both genuinely count in favor of your giving money. Now suppose that I offer you a lot of money, on the condition that you're not going to give money to famine relief *for the reason that it will make you look good*. I take it I've given you a reason not to act on that reason. In Raz's terms, I've just given you an exclusionary reason. But notice what I have *not* done: I haven't changed the fact that the excluded reason is still a reason, it still counts in favor of giving to famine relief. Nor have I added another consideration to be weighed against giving to famine relief: In fact, I haven't given you *any* reason not to give to famine relief. All I've given you was a reason not to give to famine relief *for the reason that it will make you look good*. If you go ahead and give to

famine relief for some *other* reason (like that it will alleviate terrible suffering) you can still get my money.

Sometimes it is theoretically useful to think of *combinations* of reasons. A *protected* reason, as I understand it, is such a combination of reasons: If you have a protected reason to ϕ then you have both a reason to ϕ and an exclusionary reason excluding at least some of the reasons against ϕ -ing. Perhaps, for instance, we all have a protected reason not to harm the innocent. If so, this means that we have both a reason not to harm the innocent, and a reason not to act on at least some of the reasons to harm the innocent (like, for instance, that it will help us make a buck, when it will).

Armed with exclusionary and protected reasons, we can now see how Raz hopes to accommodate the preemptiveness of authoritative directives: When a legitimate authority issues a directive " ϕ !", the subject gets a protected reason to ϕ ; that is, she gets both a reason to ϕ , and a reason not to act on at least some of the reasons not to ϕ (like, for instance, the reasons that the authority itself considered, or perhaps should have considered, in determining whether to issue the relevant directive). When you tell your 7-year-old son to go to his room, he gets both a reason to go to his room, and a reason not to stay in the living room for the reason that his favorite show is on (not *all* of his countervailing reasons need be excluded; perhaps, for instance, he can with reason stay in the living room for the reason that his room is flooded, unbeknownst to you).

How is any of this related to duties, though? Raz thinks of protected reasons as the right way to understand duties. Duties just are, Raz seems to think (with minor qualifications) protected reasons³². And notice how neatly everything fits together if he is right: For we then have an account of the preemptiveness of authority that ties it to the plausible hypothesis I've been making much of: the thought that authority consists in the power to robustly give duties. Authoritative directives generate duties, that is they generate protected reasons; and this explains their preemptive nature as well.

³² The main qualification I am aware of is that duties have to be *categorical* protected reasons (e.g. Raz 2010, 291). While this qualification is certainly important, I don't think it is important for my purposes here.

The next question to ask, of course, is when is it that such protected reasons are created, and furthermore are created by the mere say-so of another? Even if Raz's theory of duties in terms of protected reasons is right, and even if this theory can nicely explain how the directives of a justified authority – if there is such a thing – can be preemptive in the needed way, how can it be that an authority thus understood is ever justified? How can one generate a reason for another not to act on what throughout continue to be good reasons? It is here that the Normal Justification Thesis is supposed to do work. The compelling, intuitive thought is that if a subject will conform better to the reasons already applying to him, independently of the authority's directive, by obeying the authority than if he tried to act for those reasons directly (in a way not mediated by the authority's directive), then when the authority issues a directive the subject not only gets a reason to act accordingly, but also a reason not to act on at least some of the reasons not so to act³³. If so, then given the Razian understanding of duty, we also have a *duty* to act accordingly – for a protected reason of the kind just spelled out is just what such duty consists in. And we also have a ready account of the preemptiveness involved in authority – for the subject is not just given a reason for action, but also a reason not to act on some other reasons.

Now, I think that this story needs some amendments, which I am about to offer. But I think that when it comes to authority, this account is on the right track. Let me emphasize, though, that I am not at all sure that when it comes to *duties* it is on the right track. I want to remain neutral on the question whether duties should be understood as protected reasons in something like Raz's way³⁴. It will turn out that nothing of importance in what is to come will depend on this (this fact itself is of some interest, I think). Let me note, though, the following two implications of our decision how to go on duties here:

³³ For one fairly recent official statement of the NJT, see Raz (2006, 1014).

³⁴ It seems clear to me that at the very least more qualifications are needed. If, for instance, I both ask you to do something, and offer you money not to act for some reasons not to do it, I've given you both a first-order reason and an exclusionary one, but I have not given you (not necessarily, anyway) a duty. Perhaps some such examples can be dealt with by the categoricalness requirement, but I don't think it's likely that all can. Perhaps Raz can insist that a duty is not just the *conjunction* of a first order reason and an exclusionary one, but that what's necessary for a duty is that *one and the same thing* be both the first-order and the exclusionary reason. Or perhaps we need for an account of duties not just the first- and second-order reasons, but also the right relation between the two – perhaps, for instance, we want the exclusionary reason to be of the right kind.

First, if the Razian story of duties is right, the duty-out-duty-in principle should be rejected. This is so because on this picture duties are constructed from other normative material, in particular, from the special combination of first- and second-order reasons present in protected reasons. Second, if the Razian story of duties is wrong, but if his authority-story is roughly on the right track, then we should reject the plausible idea that authority is the Hohfeldian power to robustly give duties, for on some occasions authority will consist in the power to robustly give protected reasons which need not (we are now assuming) amount to duties. So the terminological point – whether we are willing to call (categorical) protected reasons *duties* – has wider (but still, I would say, terminological) implications. But not, I think, implications we should worry about here.

Razian exclusion is really an instance of a much wider phenomenon. Suppose, for instance, that I am trying to decide which desert to go for at the end of a social dinner. Clearly, caloric intake is a relevant consideration. That one desert is much more fattening than another surely counts (for me) against choosing it. But suppose that I also know that if I think about caloric intake at the dinner table, I will become a nuisance to my companions, and ruin the evening. And I have reasons not to let that happen. Now, this doesn't take away the normative significance of caloric intake, of course. Nor does this constitute an exclusionary reason in Raz's sense – there is no reason here not to *act* on the relevant reason (namely, that the first desert is much more fattening). Still, something very close is going on. What I have here is a reason *not to consider* another reason, not to think about it, perhaps not to deliberate on it³⁵.

Or suppose I need to make up my mind quickly on some matter. I have to decide between two options, and waiting for longer than a minute without any decision will be worse than making either decision. In such cases, all the first-order reasons for and against each option are intact; no exclusionary

³⁵ Raz is explicit about the distinction between reasons not to act on a reason and reasons not to consider a reason (when only the former qualifies as an exclusionary reason) when he criticizes Darwall for not consistently making this distinction. See Raz (2010, 299, fn 21).

reasons apply; and not even a reason not-to-consider is present (at least, as it were, at t_0). But what is present is a reason *not to deliberate on the matter for too long*.

Now think of combinations of reasons. Suppose you have, for instance, both a reason to do something, and a reason *not to consider certain other reasons* not to do it. Or suppose you have both a reason to do something, and a reason *not to deliberate for too long* about whether to do it. In such cases you don't have a protected reason in Raz's precise sense, for – to repeat – protected reasons partly consist of exclusionary reasons, and exclusionary reasons are reasons not to *act* for some reasons, not reasons not to *consider* some reasons, or reasons not to *deliberate* in some ways. But regardless of Raz's terminology, in the cases just described we do seem to have something very similar to protected reasons. Just to keep our terminology clean we can call these *quasi-protected reasons*. Quasi-protected reasons come in different kinds: They include (except for the relevant first-order reasons) reasons not to consider other reasons, reasons not to deliberate in some ways on some reasons³⁶, and perhaps also Razian exclusionary reasons (on this terminology, then, protected reasons are a particular instance of quasi-protected reasons). And there may be other kinds of quasi-protected reasons as well.

Now suppose that some person A has the Hohfeldian power to robustly give another person B quasi-protected reasons, but not quite protected reasons. Suppose, for instance, that I tell you to ϕ with the intentions characteristic of robust reason-giving; that I succeed in giving you a reason to ϕ ; and that though I do not succeed in also giving you a reason not to *act* on a certain class of reasons not-to- ϕ , I do succeed in giving you a reason not to *consider* certain reasons not-to- ϕ , or perhaps a reason not to

³⁶ Notice that this way we can have as a part of a quasi-protected reason not just a reason not to consider certain reasons *against* the relevant recommended action, but also a reason not to consider or deliberate on certain reasons *for* the relevant recommended action. This seems to me like a desirable result – in the case of authority, for instance, sometimes we show disrespect to the (purported) authority not just by considering reasons to act contrary to the authoritative directive, but also by considering other reasons (other, that is, than the fact that the authority said so) for acting *in the way the authority directed*. The point is not that by considering such reasons one *disobeys* the authority (standardly, as Raz (1986, 39) insists, what's required for conformity is just action, not anything about which reasons the subject considers); rather, the point is that considering such reasons – even if one does not end up disobeying – is not consistent with the appropriate attitude of a subject who does obey. On my account, of course, this fact is mirrored in the special intentions that are constitutive of robust reason-giving.

deliberate from scratch about whether or not to ϕ , or perhaps not to deliberate about ϕ -ing in a certain way. If we fill in the details correctly then it seems to me what we have here is a case of authority.

Quasi-exclusion and quasi-protection suffice, it seems to me, to capture the phenomenology of preemption that is plausibly considered conceptually necessary for authority³⁷.

And this is a good thing, too. For it is not at all clear that authorities – the things we often think, pre-theoretically, that possess legitimate authority – do succeed in manufacturing genuinely protected (as opposed to more generally quasi-protected) reasons. The best way to see this is to consider some of the "most common reasons capable of establishing the legitimacy of an authority" that Raz himself brings in *The Morality of Freedom*³⁸ (1986, 75): One such reason is if deciding for oneself causes anxiety or in some other way involves costs that relying on authority will help one avoid. But of course, in such a case, the relevant costs are costs *of the deliberation process*, or of considering the relevant reasons, or some such. So in this case what we seem to have is not a reason not to act on the relevant first-order reasons, but merely a reason not to consider them or deliberate on them. We have, in other words, quasi-protected rather than protected reasons³⁹. Another example of a way purportedly establishing the

³⁷ Raz (2010, 298) says that "there is no need to saddle the account of authority with a commitment to that way of explaining the preemptiveness of authoritative directives." Perhaps he too acknowledges here that a wider phenomenon than protected reason can accommodate the phenomenology of authority?

³⁸ Raz (2010, 298) has recently officially "crowned" the discussion of authority in *The Morality of Freedom* his most comprehensive discussion of these matters.

Raz mentions five such cases. Of the ones not discussed in the text, one is that of greater wisdom or expertise. But Raz's current position (2010, 301) is that far from being a paradigmatic way of establishing authority, expertise can (almost) *never* ground authority. I cannot discuss the details here, but let me stress that I think the current Raz is right here, and expertise cannot ground protected (or even quasi-protected) practical reasons – all it can ground is epistemic reasons to believe that there are practical reasons. So the expertise case is not an example of an authority succeeding in generating protected reasons. Indeed, the account of authority I develop in this paper nicely explains *why* authority cannot be grounded in expertise – once the information from the expert is conveyed to the non-expert, whether the expert intends to give the non-expert a reason for action immediately and completely drops out of the normative picture. So the conditions I specify for authority – even the conditions I specify for robust reason-giving in general – cannot be satisfied in (at least standard) expertise-based cases. The other of the five cases not discussed in the text is when the authority "is in a better position to achieve ... what the individual has reason to but is in no position to achieve." To the extent that I understand this case, I think it's one where sometimes – but not always – genuine exclusion will be relevant.

³⁹ In fairness to Raz, already in *The Morality of Freedom* he was ambiguous about this example, labeling it "a borderline case between normal and deviant justification". Perhaps he was already then sensitive to – though not

legitimacy of an authority is when it has a steadier will or is less likely to yield to temptation than the subject. But suppose that you know that you are likely to be weak-willed when it's time to choose a desert. I don't think this gives you a reason not to *act* on the reason that the soufflé looks great. It only gives you a reason not to *think about* which desert looks better, or perhaps a reason to just let me order for you and not think about this matter for yourself. If so, here too we have quasi-exclusion rather than exclusion. Another kind of case Raz mentions is a case where direct action on first-order reasons will be self-defeating, and relying on authority a good indirect strategy of conforming to reasons. But here too – depending on the details – often the self-defeat occurs not on the level of acting for the relevant reason but on the level of considering it (say, thinking in terms of "Are we having fun yet?" as a strategy for making sure one is having fun). So here too, often it's quasi-exclusion, not exclusion that does the work⁴⁰.

Depending on wider issues in the theory of practical reasoning there may be important relations between reasons not to act for certain reasons, reasons not to consider certain reasons in deliberation, and some other related phenomena that can manifest themselves in quasi-protected reasons. If it's true, for instance, that one can't act for a reason unless one considers it in deliberation, then a reason not to consider a reason in deliberation is *very* close to a reason not to act on it. This is just an example – and the antecedent of the conditional in the previous sentence is implausible. But this example can still show how the relation between exclusionary and some other quasi-exclusionary reasons can be determined by larger issues in practical reasoning. In fact, there is some reason to believe that here too Raz should have been discussing quasi-protected reasons more generally, and not just protected reasons. This is so, because he invokes (2006, 1025-7) a knowability requirement as a necessary

explicit about – the distinction between protected and quasi-protected reasons. In private correspondence Raz has acknowledged (if I understood him correctly) that this example is not an example of genuine exclusion.

⁴⁰ Joseph Raz suggested to me – as a general point, and perhaps also as a way of saving some of his examples – that sometimes quasi-protected reasons may then give rise to protected reasons in the more narrow sense. Even if this is so, though, it is the quasi-protected reasons that do the ultimate explanatory work here.

condition for the legitimacy of an authority: Roughly speaking, an authority can be legitimate only if its subjects can know that they would better conform to reasons by obeying the authority's directives. Raz then asks what degree of epistemic effort is being tacitly referred to by the "can know" locution, and he gives a *practical* answer – roughly, the subjects should with reasonable investment of resources be able to find out that they would better conform to reasons by obeying. But now that we're talking about reasonable investment of resources, we're *very* close to talk of reasons to consider and not to consider certain reasons, rather than to talk of genuinely exclusionary reasons in Raz's terms. In this way too, then, talk of quasi-protected reasons may be better suited to achieve what even Raz is after⁴¹.

My point is not that authorities can *never* issue genuinely protected (rather than merely quasi-protected) reasons. For one thing, I only considered (some of) the examples Raz brings, but it's quite possible that there are others. And even with regard to just these examples, I have not shown that they *never* involve genuine exclusion and protection. For instance, the self-defeat rationale may sometimes involve genuine exclusion – if, that is, we can convincingly fill in the details of a case where what leads to self-defeat is not *thinking* in a certain way or *deliberating* in a certain way or even acting in a certain way but rather *acting for a certain reason*. My point is more modest, then: It is that even though authorities sometimes create protected reasons, this is not what is distinctive about them. What *is* distinctive is that they create quasi-protected reasons – sometimes they are also genuinely protected in Raz's terms, but sometimes they are not. The distinction that does normative and explanatory work in the theory of authority is the one between reasons that are and those that are not quasi-protected. The distinction

⁴¹Paying attention to quasi-protected (rather than merely protected) reasons can be theoretically productive also in another way: By tying the discussion here to a discussion often conducted in the context of critical evaluations of utilitarianism. For one common strategy employed by utilitarians to deflate some self-defeat worries about their utilitarianism is to insist that the principle of utility is, on their view, a criterion of correctness; this does not imply that utilitarians should recommend that agents always deliberate by considering the principle of utility – they need not justify the principle of utility as a decision procedure. By employing this strategy, utilitarians insist, I think, that even though there are reasons to maximize utility, often there are reasons *not to consider* those reasons.

among the latter between those that are protected in Raz's narrower sense and those that aren't is of little significance for understanding authority.

But what happened to the relation between authority and duties? If Raz is right, duties as well as authorities are to be understood in terms of protected reasons. And then we get the relation between authorities and duties for free. I suggested, though, that authoritative preemption is wider, and includes also cases of creating quasi-protected reasons. Am I suggesting that such quasi-protected reasons suffice for duties? Again, I want to remain neutral on this question. What I am insisting on here is that the robust creation of quasi-protected reasons can suffice – depending on the details – for legitimate authority. On a sufficiently wide understanding of duty, this thesis is consistent with the thought that authority – always and everywhere – consists in the Hohfeldian power to give duties robustly. On a narrower understanding of duties, this thought has to be rejected, as sometimes authority consists in the power to robustly create quasi-protected reasons that do not amount to duties (and then, the wording of (i)-(iii) would have to be changed accordingly). Not much that is of importance here hangs, it seems to me, on the decision how liberal to be with one's use of the word "duty".

And so we get that authority consists in the Hohfeldian power to create duties or something close enough to duties (some suitably delineated quasi-protected reasons). An authority may be an authority in virtue of there being an independent duty to obey it; or, it may be an authority in virtue of there being some other normative structure in the background that allows the authority to create with its directives suitable quasi-protected reasons, thus accommodating the phenomenon of preemption we set out to explain⁴².

⁴² Some people think that there may be such a thing as *divine* authority; at least, we do not want the concept of authority to rule this out as an impossibility. Can the account in the text accommodate divine authority? (I thank Bas van der Vossen for pressing me on this.) Depending on the relevant theological details, the account in the text may apply straightforwardly. But it may not – in particular, on some theologies divine authority should be understood differently, perhaps according to the Constitutive Model I (in "Giving Practical Reasons") borrow from Schroeder (2005). According to this model, what explains the fact that when God orders that you ϕ you have a duty to ϕ , is that that's just what having a duty consists in, having a duty to ϕ consists in having been ordered by

We can return now to the worry about the explanatory impotence of thoughts about legitimate authorities if authorities are legitimate (when they are) in virtue of a prior duty to obey them. In Section 5 above I noted that even if for *some* cases the duty to obey grounds the authority and not the other way around, this does not imply that this is so for all relevant cases. And now we have – following Raz's NJT, and with the friendly amendment from talk of protected reasons to talk of quasi-protected reasons – a description of a class of cases where there is no prior duty to obey the authority; the authority is legitimate in virtue of some other thing (its power to robustly give suitable quasi-protected reasons); and the duty (or something very close to a duty) to obey is then explained by reference to the legitimacy of the authority. Indeed, the account here is appealingly unified: It gives a general account of authority as the power to robustly give duties and duty-like things, entailing that – and explaining why – sometimes an authority will be legitimate because of a prior duty to obey, and sometimes there will be a duty to obey because of a prior legitimate authority.

7. Debunking the Right to Rule

But you may think – and it has recently been forcefully argued against Raz – that all of this misses an important part of what authority is all about. Where an authority is present, the thought goes, we do not just have a duty to obey (be it explanatorily prior *or* posterior to the legitimacy of the authority itself). Authorities are *entitled* to be obeyed, they have, as it is often said, a *right to rule*. If your 7-year-old disobeys you, he doesn't merely stand in violation of a relevant duty. In a sense, he offended *against you*. In the terms Stephen Darwall uses in this and many other contexts (2006; 2009; 2010), there seems to be something essentially second-personal about what is going on, and an account of authority had

God to ϕ . You may think that this introduces some discontinuity in my theory of authority, so that divine authority and human authority come too far apart. But I do not think that this is so. In both cases, authority consists (roughly speaking) of the Hohfeldian power to robustly give duties. It's just that God is special in that He can give such duties robustly without a pre-existing duty, and without the analogue of the NJT holding (in a non-trivial way).

better be able to capture this. For instance, legitimate authorities are often entitled to hold disobeying subjects accountable in some way. A mere violation of a duty cannot accommodate this phenomenon, and the mere creation of protected or quasi-protected reasons certainly can't.

Now, when such thoughts are put forward⁴³ there is typically a lot that is going on, and so once again we need to draw distinctions. I am going to abstract the discussion here completely from the more general Darwall-esque second-personal outlook. This is so because I have discussed it – and the relation between it and robust reason-giving – elsewhere⁴⁴, and because if there is here a general objection to a general strategy of accounting for authority, it had better be possible to state the objection without relying on highly controversial, highly abstract, almost idiosyncratic philosophical theories. Rather, the point must be that there is something clear in the phenomenology of authority that is being missed here. What is it, then? I think that at the crux of the objection lies the thought that the duty to obey is *directional*. When there's a legitimate authority, its subjects are not just under a duty to obey, they *owe this duty to the authority itself*. Given the general close connections between rights and directional duties – perhaps they are just equivalent as Hohfeld (1917) thought, or perhaps there's a close justificatory rather than logical relation between them⁴⁵, or perhaps they are in some other way closely related – this directionality lies at the heart of thoughts about the right to rule, or being entitled to obedience.

But now consider an arbitrator⁴⁶. Arbitrators become arbitrators (or anyway, the one in my example did) in virtue of an agreement between the parties to the arbitration process. So suppose that our arbitrator A has issued a decision in the arbitration process between B and C. A has found in favor of

⁴³ Darwall (2009; 2010), Hershovitz (2011).

⁴⁴ "Giving Practical Reasons" (2011).

⁴⁵ Raz (1986, 170-1).

⁴⁶ Viehoff (2011) focuses attention on arbitration in a closely related context, but for very different purposes than mine. Interestingly, already in *The Morality of Freedom* (1986, 41) Raz discusses the arbitrator example, explicitly saying it's an authority. But he doesn't seem to see that this example – as I am about to argue in the case – is a counterexample to the right-to-rule.

B, so that C has to pay. And suppose that all is well with the case – that is, the parties fully and rationally consented to the arbitration process, A decided with reason and wisdom, the decision was well within his power (as agreed on by B and C), etc. Nevertheless, suppose that C refuses to pay. Pretheoretically, it seems to me, we want to say about this case that it's a case of authority (don't arbitrators sometimes possess legitimate authority?). And it is a case in which a duty has been generated, for C now stands in violation of his duty to follow A's decision (and pay up). But what about the directionality of this duty? *To whom* does C owe the duty of obedience here? It seems clear that he owes this duty *to B*, the other party to the arbitration agreement, and not to A, the arbitrator. He may owe A other duties (like perhaps the duty to pay her for her arbitration services), but this is a different matter. If all of this is right – in particular, if A has genuine authority, but the duty to obey is not owed to A – then we have here a counterexample to the claim that it is necessary for authority that the duty to obey be owed to the authority.

This counterexample seems entirely conclusive to me. And it seems like an instance of a more general pattern of counterexamples here, some of which perhaps more obviously politically relevant. The thought, for instance, that an authoritative democratic legislature is the one *to whom* obedience is owed seems very odd, and indeed antagonistic to the democratic spirit – it's *your fellow citizens* to whom you owe your obedience to the democratic legislature (if indeed you owe such a duty), isn't it⁴⁷? My point is not, of course, that obedience can *never* be owed to the authority. Especially given the role I've allowed for a prior duty to obey in accounting for authority, we can be very liberal here – sometimes, the prior duty to obey in virtue of which an authority emerges may be owed to the authority. The crucial point, though, is that this is not always and necessarily the case. So the directionality point Darwall and others emphasize just fails⁴⁸.

⁴⁷ Herschovitz (2011) is aware of the problem, but it's not clear to me what his solution is.

⁴⁸ Marmor (manuscript) makes a similar point, though he argues for it differently.

Indeed, the thought that authorities – certainly, political authorities – have anything worth calling a right to rule seems to be a remnant from days long gone. *No one* has a general right to rule, certainly politically⁴⁹. I concede that there is something of some intuitive force somewhere in the vicinity here, and I proceed to offer a debunking explanation of it in what follows. But we should be clear about there being nothing at all intuitive in the political sphere about a right to rule.

I think it is one of the advantages of Raz's service conception of authority – indeed, perhaps one of its original *points* – that it only allows authorities to be justified in terms of the service they render the subjects. It is, as it were, all about helping the subjects (conform to reasons), not about the authority itself. Precisely for this reason, though, I find it perplexing to find Raz himself routinely talking in terms of the authority's right to rule. And while in earlier texts one may have thought that this was just a figure of speech, in more recent work⁵⁰ Raz adds even more perplexing details. Raz seems to insist (and is read by others⁵¹ as insisting) that authority consists in a right to rule; Raz then plugs in his interest theory of rights, so that (even less plausibly!) authority must somehow be grounded in the authority's *interests*; he then locates the relevant interest – the interest of the authority to govern well, or to be a good authority; and argues that the right to rule is grounded in that interest. Now, several things may be wrong with this way of proceeding, but I want to focus on just one. Even assuming that authorities have an interest in governing well or in being good authorities (it seems to me plausible that they often do, but highly implausible that they always and necessarily do), the thought that their status as legitimate authorities is somehow *grounded* in this interest of theirs (or in any other interest *of theirs*) seems both terribly implausible, and in direct tension with the powerful insights underlying the service conception

⁴⁹ I am not at all sure I understand Edmundson's (2010) discussion of authority as intrinsically valuable and giving rise to intrinsic reasons to obey. To the extent that I do, though, let me note that I would respond to it in a way analogous to the point in the text here – authority is *never* intrinsically of value.

⁵⁰ See here mostly Hershovitz (2011, 8-9), and the references there, mostly to Raz's manuscript "The Possibility of Partiality".

⁵¹ Again see Hershovitz (2011). See also Sherman (2010). In conversation, Raz rejected this reading. See below.

of authority. Indeed, in his most recent text on the matter, Raz seems to insist – in replying to Darwall – on precisely such a tension⁵².

If I am right, then, thoughts about authorities' right to rule – as a conceptually necessary condition for their legitimacy – should, unless interpreted *very weakly*⁵³, be debunked. But I do not want to deny that there is some intuitive force underlying such thoughts. How can this be explained? In the political context, perhaps historical explanations can play some debunking role here – talk of a right to rule, in the political context, probably originates with attempts to vindicate the authority of rulers we would not today count as legitimate, and perhaps it is time to reject not just those regimes but also the understanding of authority that seems to have come with them⁵⁴. But can more be done here to explain away intuitions about a right to rule?

⁵² "I am more used to the idea that those in authority are accountable to their subjects than to the thought that their subjects are accountable to them." (2010, 299) And in conversation Raz confirmed that he does not think authority ever *consists in or is grounded in* the right to rule. Rather, it's just that authorities (or anyway some authorities, mostly authoritative institutions) often *have* a right to rule. I need not reject this weaker claim. Leslie Green suggested to me that thoughts about the right to rule become more plausible if we bear in mind that while such rights must be grounded in authorities' interests, they need not be *ultimately* so grounded, in that the authorities interests may be important here merely instrumentally, and indeed in virtue of other people's interests. (See here, for instance, Raz's (1986, 179) discussion of the journalist's right not to disclose her sources, which is grounded in her interest to collect information, which is itself of value mostly because of the *public's* (not the *journalist's*) interest in the flow of information.) I agree that this would make talk of a right to rule *more* plausible. But it wouldn't make it plausible. Presumably, when a journalist's right not to disclose her sources is violated, the journalist is in some way especially entitled to complain. Not so, it seems to me, for my arbitrator, or indeed for authorities in general.

⁵³ Several weak interpretations of talk of the right to rule can now be found in the literature. They include understanding it as merely designating a Hohfeldian power to generate duties (Perry (manuscript)); as a Hohfeldian power to generate *legal* duties (Applbaum 2010); and maybe even as a Hohfeldian privilege (so that, say, in many circumstances many of us have a duty not to interfere with the actions of others, but someone who is a legitimate authority does not have such a duty; I thank Ben Saunders for this suggestion). When I debunk the right to rule in the text, what I debunk is the Hohfeldian *claim* right to rule. I accept, of course, the power to generate duties; I don't think that Applbaum power to generate *legal* duties suffices for the kind of authority that may be normatively interesting, but I acknowledge that the power to generate legal duties is a coherent phenomenon worth thinking about in this context (and that it's not vulnerable to what I say in the text about the right to rule); and I have no problem with a Hohfeldian privilege to rule, though I think that – as the literature addressed throughout this paper makes clear (see here also Simmons (2001, 130)) – it doesn't capture the phenomenon of authority we're all interested in.

⁵⁴ Alternatively, if you think that this feature has become an essential part of our concept of authority, then what I am suggesting here is that we go error-theoretic on authority. But I find it hard to get all excited about talk of error theory here, because I am not at all sure that talk of a right to rule is a part of our everyday concept of authority (as opposed to merely being a part of philosophical dogma here), and because what I am offering as authority

Think about our arbitrator again. I was suggesting that the duty of the parties to follow her decision is not owed to the arbitrator (but to each other). So when C violates this duty, he does not violate a duty he owes A. Still, is A not in some way entitled to complain? If A feels that she too is wronged by C's violation, must we say that she is mistaken? True, as we sometimes say, it's not (primarily) about her, it's about B's right against C. Nevertheless, I think that if A feels offended, wronged, somehow more personally involved than merely as a service-provider, this need not be irrational on her part. The thing to note here, I think, is that while C's duty to follow the arbitrator's ruling is not owed to the arbitrator, and so when it is violated the party wronged is not the arbitrator, still some *other* things that are going on here may involve wronging the arbitrator. For in some circumstances, failing to obey an arbitrator's ruling may have an expressive element to it – it may involve something like saying that the arbitrator is not a good arbitrator, or did not in good faith attempt to reach a fair decision, or some such. In saying such things, C may very well be wronging the arbitrator, indeed perhaps even violating a duty he owes to the arbitrator. And so A is entitled to feel offended, and perhaps (depending on other details) A also has the standing to complain or demand reparation or apology or some such. And this gets us even closer to the directionality talk of right to rule was meant to capture: Yes, often disobeying legitimate authority involves a violation of a duty owed to the authority – namely, the duty not to (falsely) claim that it is a bad authority. But the duty relevant here is not the one that is arguably necessary – as a conceptual matter – for authority. And it is not universally present, certainly not to the same degree – whether and the extent to which it is will depend on the details of the specific authority, pragmatics of conversation, whether disobedience is understood as reflecting badly on the authority, etc. We can get further evidence for the plausibility of this explanation by

gives us so much of what we pretheoretically want, that even if it doesn't give us everything, it's hard to view it as an error theory. In this case, then, I side with David Lewis (1989, 93), who ridicules suggestions to go error-theoretic on every-day ascriptions of simultaneity because of modern physics' problems with simultaneity: "Shock horror: no such thing as simultaneity! Nobody ever whistled while he worked!"

thinking about cases in which a duty to obey is not plausibly thought of as owed to the authority but to others (as in the arbitrator case) or perhaps to no one at all; and where violation will not be plausibly construed as saying that the authority is a bad authority. In such cases, it seems to me, thoughts about right to rule, standing to hold accountable, the authority's entitlement to be obeyed, and the directionality of the duty to obey lose all of their intuitive support. And even in those cases in which violation will be understood as saying (roughly) that the authority is a bad authority, and in which the authority therefore has an interest not to be disobeyed, nothing about this interest will *ground* the authority.

The understanding of authority as the Hohfeldian power to robustly give duties (or duty-like things), backed by the combined normative story of prior duties to obey and the friendly-amended NJT, is indeed conceptually insensitive to the directionality of the duty to obey. But we have examples showing that this is as it should be, and we have a plausible explanation of why we would have intuitions about such directionality even if they are misleading (namely, because there often are *other* duties with the relevant directionality in the vicinity, like the duty not to falsely say that the authority is a bad authority). So we do not have here a convincing objection to this understanding of authority⁵⁵.

8. A quick Word about Epistemic Authority

What follows from all of this to *epistemic* authority? I think it follows – as it should – that there is nothing like epistemic authority, if this is understood as something that is supposed to exactly parallel practical authority. Let me explain.

⁵⁵ Japa Pallikkathayil suggested to me that while *some* cases of authority – like the arbitrator's – do not involve a right to rule, others do (the parent case was the case she mentioned). I agree that there may be cases of authority where a right to rule is also present (though I am not sure about the parent case); it's just that the right to rule is not a part of what it is for the authority to be an authority. Also, I can accommodate the intuitive feeling – if it is that – that there are distinctions here between the arbitrator and the parent case using the explanation in the text: Surely, it is a much greater offense against me to say that I'm a bad parent than to say that I'm a bad arbitrator.

Epistemic reasons – standard reasons for belief – can certainly be given. For one thing, and rather obviously, they can be epistemically given, as when I indicate to you some piece of evidence you hadn't noticed earlier. And it seems likely that they can be triggered as well – perhaps, say, by conducting some new experiments. But I think that epistemic reasons cannot be *robustly* given. This just follows from a rather uncompromising insistence on the distinction between practical and epistemic reasons (for instance, for beliefs). This follows from my theory of robust reason-giving, as I am about to show, but notice first that there is also something very intuitive about this – If I ask not just that you read my paper, but that you believe it's very good, there seems to be something confused about my request, and it doesn't seem that you can accommodate it (by believing that the paper is good for the reason that I asked you to) in the same way you can accommodate the more standard request (by reading the paper because I asked you to). Of course, someone who is confused about these things may *try* to give epistemic reasons robustly, as is shown by the example of the confused request that you believe my paper is very good. But such attempts cannot succeed, because the normative success condition is not met: The fact that someone forms (i)-(ii)-(iii)-like intentions that you believe that p is not a reason for you to believe that p⁵⁶.

This, as I just said, is true of robust reason-giving in general, and so it's true for the case of authority in particular. One cannot succeed in issuing a command that someone believe something⁵⁷. This seems clear pre-theoretically, and it is also the conclusion that follows from my theory of authority: One can *attempt* to issue such a command (though, of course, such an attempt would already be confused and highly non-standard), but one cannot succeed, because the normative success condition is not met – there is no epistemic reason to believe that can be triggered by such a command.

⁵⁶ There's more on this – including what this implies with regard to the assurance view of testimony – in "Giving Practical Reasons".

⁵⁷ Raz (2006, 1034) makes this point, though in a somewhat different context.

But still, we do sometimes speak of epistemic authority. Must I reject all such talk as confused? If so, isn't this a liability? Don't we want a theory of authority that gives a somewhat unified account of practical and epistemic authority⁵⁸? Well, as the previous paragraphs show, we do not want an account of epistemic and practical authority to be *too* unified. But the account suggested here is, I think, unified to just the right extent. For the idea of protected and indeed quasi-protected reasons can be easily applied in the epistemic case. And where we have quasi-protected reasons, we can have preemption. Sometimes, one would believe better (in a way that is better in accordance with the evidence, perhaps, or that is more reliable) by relying on the say-so of another than by trying to respond directly to the first-order epistemic reasons. This, I take it, is what goes on in the case of expertise. Experts succeed in giving others quasi-protected reasons for belief. But notice that they don't give these reasons *robustly* – they merely trigger them.

If you want, you can say that experts have epistemic authority. Here as anywhere else, nothing depends on terminology. But we should be clear about what experts do and what they do not have in common with practical authorities: They share the ability to give quasi-protected reasons, and so to generate the phenomenology of preemption. They differ in that only in the practical case can these reasons be given robustly (so that epistemic reasons can't be given by requesting, ordering, promising, etc.).

9. A Quick Word about the (Ir)Relevance of Practices

Scott Hershovitz (2011) has recently argued that authority cannot be made sense of independently of certain social practices. He writes (2011, 11): "Authority is a feature of roles embedded in practices. To justify authority, we need to justify the practices in which roles of authority are embedded."⁵⁹ On this

⁵⁸ I thank Bas van der Vossen and Ofer Malcai for pressing me on this point.

⁵⁹ For a closely related view, see also Marmor (manuscript). The critical points that follow in the text apply, I think, to Marmor's version as well.

view, as I understand it, it is conceptually impossible that there be legitimate authority independently of some relevant social practices. Social practices, on such a view, are a part of what constitutes authority.

Obviously, much more can and should be said about the relation between authorities and practices (think here, for instance, about the volumes of literature on the relation between promises and practices). But let me settle for just the following two points.

First, though I agree with Hershovitz that there is a close relation between authority and practices, I do not think the relation is *as* close as he makes it. For on any non-question-begging understanding of what a practice consists in, the thought that it's conceptually impossible for there to be legitimate authority without a relevant social practice seems implausible to me. The only way to render it plausible, it seems to me, is to understand "practice" in a way that guarantees that this is so. But of course, we are not interested in linguistic legislation of this kind.

Second, even if Hershovitz is right and there cannot be authority without an underlying social practice, we may still ask for an explanation of this surprising fact. Why is it, in other words, that authority cannot be justified independently of practices? And in answering this question, I think we will need to rely on a theory of authority – what it is, and how it can be justified – that *explains* rather than presupposes the relevance of practices. In other words, given that the relation between authorities and practices is highly unlikely to be brute (even if it is necessary), we need a deeper understanding of authority, one that does not rely on practices.

On both these points, the account suggested in this paper does better. For arguably, practices are relevant to authorities because they often determine when the normative success condition is met – when, that is, either there is a pre-existing duty to obey, or the relevant authority can generate quasi-protected reasons for action for its subjects. If it is impossible for this (disjunctive) condition to be met without an underlying social practice, then practices will end up being necessary for authority, but in a way that is explained by the deeper account of authority. And if – as seems more likely to me – the

normative success condition can be met without an underlying social practice, then the relation between authorities and practices will end up being contingent (though perhaps close nevertheless). And this too will be nicely explained by the account of authority suggested here.

10. So: When Is an Authority Legitimate?

I have put forward here an account of authority as the Hohfeldian power to robustly give duties (or duty-like things, quasi-protected reasons). Armed with this understanding of authority, we can now return to the two problems we started with.

The first, metaphysical, problem was about the ease with which we seemingly can introduce new duties all the way out there in Plato's heaven. But having realized that robust reason-giving is a particular instance of triggering reason-giving, we now know that there's nothing more metaphysically mysterious about authorities creating new duties than about a grocer raising the price of milk thereby giving you a reason to buy less milk⁶⁰.

How about the moral problem, then? How can it be justified for one person to subject his will to the will of another? First, I argued, if there is a prior duty to obey someone, then this duty will explain why one person should here subject his will to that of another (and will ground authority in the process). I had nothing to say about when such duties exist, though. But one plausible class of cases here is where consent makes a difference. Perhaps, in some cases of consent (or perhaps promise, or commitment, or agreement, or some such) to obey, a duty to obey emerges, and with it the duty to subject one's will to another's, and the existence of authority. If, for instance, I promise to obey you (within a certain time

⁶⁰ Raz (2006, 1020) says that the reasons we should adopt his service conception is that it solves the metaphysical and the moral problems regarding authority. But with regard to the metaphysical problem, this is not quite so. Nothing in the particular details of the service conception at all helps with that problem. All that does is noticing that we have here an instance of triggering reason-giving, something that pretty much all of those rejecting the service conception can also accept. For Raz's acknowledgment that authority is a particular instance of triggering reason-giving (though not using these terms, of course), see Raz (1986, 84).

frame, on a certain restricted class of issues, assuming none of a class of exceptions applies, etc.) then it seems like I should obey you, that is, I should subject my will to yours. And it doesn't seem like in such cases much of the initial moral mystery about such authority remains (unless, that is, a mystery about *promises* remains).

But cases where there is a prior duty to obey do not necessarily exhaust the cases where we can answer the moral problem of authority. Following Raz on protected reasons, I argued that it's quite possible that sometimes even if there is no prior duty to obey, still we can robustly give each other reasons that are quasi-protected, in that they are accompanied by (also given) reasons not to act on other reasons, or not to consider certain reasons, or not to deliberate on the matter in a certain way. In such cases, I argued, enough of the phenomenology of authority is present for us to justifiably view them as cases where authority is present. And so we have – in outline, at least – two related ways in which the moral problem can be satisfactorily dealt with.

But all of that – viewed from the perspective of political philosophy – has been merely preliminary work. For I have said close to nothing on when authorities are actually justified, much less have I said anything more specific about *political* authority. Under what real-life conditions are which persons and multi-person bodies legitimately authoritative, on what matters, and within what constraints?

I hope to engage some of these questions in detail in the future. I certainly do not want to pretend that in this paper I already have. But nor do I think that what's been achieved in this paper is irrelevant for these more practically-pressing matters. If the account of authority presented here is right, we know how in outline we should proceed on these matters. We should check to see when there are prior duties to obey. And we should check to see when we can succeed in robustly giving each other quasi-protected reasons – when, in other words, a person or a person-like thing can form the intentions

constitutive of robust reason-giving and indeed of robustly giving duties, thereby either successfully triggering a pre-existing duty or successfully creating a quasi-protected reason.

I think this strategy shows promise in dealing with the more-or-less paradigmatic cases. When you tell your 7-year-old to go to his room, you form the relevant intentions, and you succeed in thereby creating a duty for him to go to his room – either because he has a general duty to obey you, or because in the situation it is within your Hohfeldian powers to robustly create quasi-protected reasons (often deliberation for someone with the deliberation skills of a 7-year-old can be very costly, so often he has reason not to deliberate but rather to do as you say).

Can states possess legitimate authority? A small question this isn't, and I hope to address it elsewhere (surprised?). But let me just state that though I do not think that consent-theories can succeed here, I can easily accommodate their appeal – if consent could be secured, this would have been a way of grounding authority in a prior duty to obey, itself grounded in consent. Perhaps something analogous can be done with more sophisticated theories here – perhaps, for instance, an attempt at grounding the legitimacy of democratic states in considerations of procedural fairness can be accommodated by noticing that in such cases there is a prior duty (grounded in considerations of fairness) to follow the outputs of the fair procedures. And so on.

As for the volunteer passenger who solves crucial coordination problems – I have already said that I think she lacks authority, at least initially. I also indicated how if a practice of following her instructions is established, gradually authority may emerge. And let me note that more generally, the ability to solve coordination problems may be a paradigmatic – *the* paradigmatic, I am tempted to say – way in which authority may emerge without a prior duty to obey. And here too, it seems to me that quasi-protected reasons are better suited as an explanation than Raz's protected reasons.

The substantive work, then, by-and-large remains to be done. But I hope that the work done here can help to do that work in a more productive way. But the proof of this pudding too is in the eating. So

the value of the understanding of authority developed here will ultimately be determined to a large extent by the work it can do – *if it can* – in the more substantive discussions of authority.

References

- Arthur Isak Applbaum (2010) "Legitimacy without the Duty to Obey", *Philosophy and Public Affairs* 38, 215-239.
- Geoffrey Cupit, "How Requests (and Promises) Create Obligations", *The Philosophical Quarterly* 44 (1994), 439–455.
- Stephen Darwall, *The Second-Person Standpoint: Morality, Respect, and Accountability* (Cambridge, MA: Harvard University Press, 2006).
- (2009) "Authority and Second-Personal Reasons for Acting", in David Sobel and Steven Wall (eds.) *Reasons for Action* (Cambridge University Press, 2009), 134-154.
- (2010) "Authority and Reasons: Exclusionary and Second-Personal", *Ethics* 120 (2010), 257-278.
- William A. Edmundson, "Political Authority, Moral Powers and the Intrinsic Value of Obedience", *Oxford Journal of Legal Studies* 30 (2010), 179-191.
- David Enoch (2011), "Giving Practical Reasons", *Philosophers' Imprint* 11(4), 1-21.
- (2011) *Taking Morality Seriously: A Defense of Robust Realism* (Oxford University Press).
- (forthcoming) "Being Responsible, Taking Responsibility, and Penumbral Agency", forthcoming in Ulrike Heuer and Gerlad Lang (eds.), *Luck, Value, and Commitment: Themes from the Ethics of Bernard Williams* (Oxford University Press).
- David M. Estlund, *Democratic Authority: A Philosophical Framework* (Princeton and Oxford: Princeton University Press, 2008).
- H. L. A. Hart, "Commands and Authoritative Legal Reasons", in his *Essays on Bentham* (Oxford: Oxford University Press, 1982), 243–268.
- Scott Hershovitz, (2011) "The Role of Authority", *Philosopher's Imprint* 11(7).
- Lewis, David. 1969. *Convention*. Cambridge: Harvard University Press.
- (1989), "Dispositional Theories of Value", *Proceedings of the Aristotelian Society*, supp. Vol. 63, 113-137; reprinted in his *Papers in Ethics and Social Philosophy* (Cambridge: Cambridge University Press, 2000), 68-94.

Wesley Newcomb Hohfeld (1917), "Fundamental Legal Conceptions as Applied in Judicial Reasoning", 26 Yale Law Journal, 710.

Andrei Marmor, (manuscript), "The Dilemma of Authority".

Stephen Perry (manuscript) "Political Authority and Political Obligation".

Joseph Raz, *Practical Reason and Norms* (London: Hutchinson & Co., 1975).

(1986) *The Morality of Freedom* (Oxford: Oxford University Press, 1986).

(2006) "The Problem of Authority: Revisiting the Service Conception", 90 *Minnesota Law Review*, 1003–1044.

(2010) "On Respect, Authority, and Neutrality: A Response", *Ethics* 120 (2010), 279-301.

Stephen Schiffer, *Meaning* (Oxford and New York: Oxford University Press, 1972).

(1987) *Remnants of Meaning* (Cambridge, MA: The MIT Press, 1987).

Mark Schroeder, "Cudworth and Normative Explanations", *Journal of Ethics and Social Philosophy* 1(3) (2005), available on line at www.jesp.org

Stefan Sciaraffa, "On Content-Independent Reasons: It's Not in the Name", *Law and Philosophy* 28 (2009), 233–260.

Matthew Noah Smith, "Political Obligation and the Self", forthcoming in *Philosophy and Phenomenological Research*.

James Sherman, "Unresolved Problems in the Service Conception of Authority", *Oxford Journal of Legal Studies* 30 (2010), 419-440.

A. John Simmons (2001), *Justification and Legitimacy: Essays on Rights and Obligations* (Cambridge: Cambridge University Press).

Daniel Viehoff (2011), "Debate: Procedure and Outcome in the Justification of Authority", *The Journal of Political Philosophy* 19 (2), 248-259.