

## The Sources of Authoritative Exclusion

*(penultimate draft – please do not cite)*

It is often held that authoritative directives are intended to serve as exclusionary reasons, meaning that their intended force is that those who are subject to them should not act on or be motivated by some of their other reasons for action. I argue that the exclusion found in two common domains (practical rationality and non-instrumental moral duties) differs too much in character from authoritative exclusion to plausibly serve as its single source. I argue that a better account of authoritative exclusion can be constructed by combining those two sources, providing support for what I describe as a moralized version of Joseph Raz's normal justification thesis, which understands obedience to authority as a requirement of instrumental morality rather than instrumental rationality.

### **1 Exclusion**

A common picture of reasoning about action is that we should ordinarily be motivated by our reasons in proportion to how much their contents speak for or against performing the relevant actions.<sup>1</sup> Many philosophers who otherwise accept this picture add a nuance to it: they hold that there are some cases in which we are permitted or even required to *exclude* certain considerations from our deliberations that lead to action. There is substantial disagreement among philosophers about the finer details of how exclusion works, whether it is a unified phenomenon, or even whether it is coherent at all.<sup>2</sup> Nevertheless, to a first approximation, we can illustrate exclusion with the following simple model of deliberation and action: ordinarily in deciding whether or not to perform an action, an agent will take all of the reasons that speak for or against it, add them up, and then either do or not do the action depending on which way the balance of reasons tilts. When an agent excludes a reason from their deliberations about how to

---

<sup>1</sup> I thank David Dyzenhaus, Arthur Ripstein, Joseph Heath, Kenneth Ehrenberg, Hamish Russell, as well as two anonymous referees for *Law and Philosophy*, for thoughtful comments on various versions and antecedents of this paper.

<sup>2</sup> E.g. Stephen Perry, "Second-Order Reasons, Uncertainty, and Legal Theory", *Southern California Law Review*, 62, (1989): 913-994, and Joseph Raz, "Facing Up: A Reply", *Southern California Law Review*, 62, (1989): 1153-1235.

act, it plays no role in determining how they ultimately act, even though it still seems to speak for or against performing the relevant action.<sup>3</sup>

## 2 Candidate Sources of Authoritative Exclusion

Exclusion is particularly important to political philosophy and philosophy of law because of the central role that it is thought to play in *authority*. Many philosophers hold that authorities claim not only to give their subjects strong reasons or duties, but also to exclude some of their other reasons.<sup>4</sup> When those authorities are legitimate, their claims are valid, and those subject to them should both act in accordance with those claims and exclude some of their other reasons for action.

Although many philosophers hold that authority involves exclusion, they have found it very difficult to explain how claims to authoritative exclusion could ever be valid. There seem to be many barriers to the validity of such claims. For example, how can it be rational to do what someone says, just because they tell you to do it? And given the apparent moral importance of making decisions for ourselves, how can it be (morally) permissible to allow someone else's will to replace one's judgment? More difficulties pile on when trying to explain the validity of the extraordinarily ranging and demanding claims made by political and legal authorities. Yet

---

<sup>3</sup> Throughout this paper, I follow Joseph Raz in assuming that exclusionary reasons are motivating reasons – reasons to be or not be motivated by certain reasons. I think it is an open question whether the argument of this paper extends to exclusionary reasons when they are understood as normative reasons. For a detailed argument that exclusionary reasons are normative reasons, see N.P. Adams, “In Defense of Exclusionary Reasons”, *Philosophical Studies*, 178, (2021): 235-253.

<sup>4</sup> In *Practical Reason and Norms*, Joseph Raz introduces the concept of an exclusionary reason and analyzes norms in terms of it, while slightly later in *The Authority of Law*, he gives an explicit analysis of authority in terms of exclusionary reasons. See Joseph Raz, *Practical Reason and Norms*, (Oxford, UK: Oxford University Press, 1999) and Joseph Raz, *The Authority of Law (Second Edition)*, (Oxford, UK: Oxford University Press, 2009). Pre-dating Raz, Robert Paul Wolff also seems to understand claims to authority in terms of exclusion, though he does not use the concept explicitly. See Robert Paul Wolff, *In Defense of Anarchism*, (New York: Harper and Row, 1970).

despite these difficulties, there have been many attempts to provide such an explanation, and my first aim in this paper is to draw attention to a feature of many of these attempted explanations that is often left implicit. It is that they are *indirect*: they first identify some form of exclusion found in a source outside of political authority and then extend the principle explaining that exclusion to the harder case of authoritative directives.

Joseph Raz's normal justification thesis is arguably the first and most influential of these indirect explanations. In short, he observes that there are many aspects of practical reason, including reasoning under non-ideal conditions and decision-making, that appear to neatly lend themselves to analysis in terms of exclusionary reasons. For Raz, it was a natural jump from these cases of exclusionary reasons to the service conception, and especially the normal justification thesis within it, which understands the role of exclusion in authority as helping agents act on the reasons that already independently apply to them.<sup>5</sup>

Even though Raz in later work almost always presents the service conception of authority in isolation from his earlier observations about practical reason, it is nonetheless useful to present it as flowing from these earlier observations for two reasons. First, this way of presenting things is an important (if implicit) part of the *argument* for the service conception. Though Raz claims to present the service conception as a conceptual analysis of authority, part of why we find that conceptual analysis plausible in the first place is that the examples of exclusion he draws from practical rationality independently ring true to us.<sup>6</sup> If we already grant that sometimes we should

---

<sup>5</sup> Raz's normal justification thesis is just one part (though the most important part) of his service conception of authority.

<sup>6</sup> See, for example, Joseph Raz, "Revisiting the Service Conception", in *Between Authority and Interpretation: On the Theory of Law and Practical Reason* (Oxford: Oxford University Press, 2009): 129-134.

not act on or be motivated by our reasons (say, when we are too tired to make an effective decision), then it becomes somewhat easier to accept that we can or should do the same when confronted by authorities. Second, this presentation is helpful for identifying potential *competitors* to the normal justification thesis and ensuring that they are not talking past one another. Many apparent competitors to the normal justification thesis, like many versions of fair play theories and consent theories, are drawn from earlier debates about the duty to obey the law. Since the philosophers who advanced those theories were not clear about whether they thought exclusion was a constitutive feature of the duty to obey the law, it is difficult to tell whether their views are, or even could be, genuine competitors to the normal justification thesis.

Thus, rather than looking for those competitors by looking at existing theories of the duty to obey the law, I will instead look for them by pursuing the same strategy that I have attributed to Raz. Are there any other domains, apart from practical rationality, in which we find exclusion, that we can then try to extend to authoritative exclusion?<sup>7</sup> I will suggest that there is at least one further source of exclusion apart from practical reason: non-instrumental moral duties. The next step of my argument will compare the exclusion found in practical reason and non-instrumental moral duties with the exclusion found in authoritative directives. Is authoritative exclusion more like the exclusion found in practical rationality, or the exclusion found in moral duties, or does it differ altogether from the exclusion found in these domains? And what can these similarities and differences tell us about the prospects of each of them explaining the exclusionary force claimed by authorities, particularly legal and political authorities?

---

<sup>7</sup> For a useful discussion of the many uses to which Raz himself has put exclusion, see Michael Moore, "Requiem for a Concept: Exclusionary Reason", in Andrei Marmor, et al (eds.), *Essays in Honour of Joseph Raz* (Oxford: Oxford University Press, forthcoming).

Note one important limitation of this next step of my argument. Even if we assume that each of the source domains which I discuss genuinely does involve exclusion, and that the source domain shares the same properties as authoritative exclusion, this does not necessarily show that authoritative exclusion can be grounded in this source. On most accounts, authoritative exclusion involves not only excluding certain reasons, but also taking a command as a first-order reason to do what is commanded, regardless of its content. Thus, even if there are similarities in features between the exclusion found in one of the previous sources and authoritative exclusion, it would still take further work to show that the source in question can explain the validity of a whole claim of authority. What this strategy accomplishes is thus mostly *negative*: it can show that some grounds of exclusion do not suffice, on their own, to ground authoritative exclusion.

### **3 Source 1: Practical Rationality**

I will now say more about the two domains. The first I have already briefly discussed: the instrumental concerns of practical rationality itself. Consider Raz's example of someone who is looking to invest some of her money.<sup>8</sup> Suppose that she must decide whether to take a promising investment opportunity within a few hours, but she is presently exhausted and unable to make an effective decision. How should she take the fact that she is exhausted into consideration in her reasoning? It seems like it must play some role in her reasoning, but it is not properly speaking a property of the investment, and so not eligible to be a first-order reason. Instead, this fact is a 'second-order' reason, that is, a reason about a reason. There are many kinds of second-order reasons, but the investor's fatigue is a particular kind of second-order reason that Raz calls an

---

<sup>8</sup> Raz, *Practical Reason and Norms*, 37.

‘exclusionary’ reason. An exclusionary reason is a reason to ‘disregard’ our other reasons or not be motivated by them.<sup>9</sup> (Raz also frequently uses the idea of a ‘protected’ reason—a reason that exclude some reasons but also provides a first-order reason in favour of some action.) In the case of the investment, for example, the fatigue gives the potential investor an exclusionary reason to not be motivated by their reasons in favour of making the investment. Thus, it seems like we find a rationale for excluding reasons in the internal logic of practical reasoning itself—in at least some cases, to make it more likely that our reasoning about how to act will reach the right conclusion, we should exclude some reasons from our deliberations about how to act.

In his service conception of authority, Raz draws upon the same kind of considerations to explain why authoritative directives might be exclusionary. The centerpiece of the service conception is the *normal justification thesis (NJT)*:

The normal and primary way to establish that a person should be acknowledged to have authority over another person involves showing that the alleged subject is likely better to comply with reasons which apply to him (other than the alleged authoritative directives) if he accepts the directives of the alleged authority as authoritatively binding and tries to follow them, than if he tries to follow the reasons which apply to him directly.<sup>10</sup>

---

<sup>9</sup> Raz, *The Authority of Law*, 18. He clarifies in the postscript to the second edition of *Practical Reason and Norms* that he means that exclusionary reasons are reasons “for not acting” on certain reasons, or reasons “for not being motivated in one’s action by certain (valid) considerations”. (Raz, *Practical Reason and Norms*, 184-5)

<sup>10</sup> Joseph Raz, “Authority, Law and Morality”, *The Monist*, 68, (1985): 299.

In short, we should acknowledge that someone (or something) has legitimate authority over us if doing so would help us better conform to our underlying, background reasons, which Raz calls ‘dependent’ reasons. He thinks that this is likely to occur in five kinds of cases:

1. When the authority is “wiser.”
2. When the authority has a “steadier will less likely to be tainted by bias, weakness or impetuosity.”
3. When direct individual action is “likely to be self-defeating.”
4. When deciding for oneself “causes anxiety, exhaustion, or involves costs in time or resources...”
5. When the authority is “in a better position to achieve.. what the individual has reason to but is in no position to achieve” (here Raz seems to be referring to the role of authority in solving coordination problems).<sup>11</sup>

Thus, the service conception implies (for better or worse) a piecemeal conception of authority, on which the state may have greater or lesser authority over individuals depending on how equipped they already are to act on their own reasons without the assistance of the law.

Raz introduces one qualification to the normal justification thesis that will be important later in my argument. He suggests that authority will be ‘binding’ only if the underlying reasons are *categorical*, that is, independent of the beliefs and desires of the subject.<sup>12</sup>

---

<sup>11</sup> Raz, *The Morality of Freedom*, 75

<sup>12</sup> Raz, *The Morality of Freedom*, 64.

In another part of the service conception, the pre-emption thesis, Raz suggests that authoritative directions that meet the conditions of the normal justification thesis should be treated as exclusionary reasons:

The fact that an authority requires performance of an action is a reason for its performance which is not to be added to all other relevant reasons when assessing what to do, but should exclude and take the place of some of them.<sup>13</sup>

The pre-emption thesis states that we should treat authoritative directives as exclusionary reasons but does not supply an argument for *why* we should do so. Even if an authority satisfies the preconditions of the normal justification thesis, why should we treat its directives as exclusionary? The main argument that Raz presents is a double-counting argument: if authoritative directives depend on the reasons that subjects already have, then authoritative directives must exclude those reasons, lest the subject irrationally count them twice.<sup>14</sup> However, the double-counting argument is incomplete, since it does not explain *why* we should take those directives as reasons at all (let alone exclusionary reasons) in the first place. Raz provides a more direct answer to this question in later work where he clarifies that the exclusionary reasons created by the traffic light exist “by reasoning analogous (some would say identical) to that which establishes the existence of instrumental reasons: you have reason to do A, doing B (walking to the station, obeying the authority) will facilitate doing A, therefore you have reason to do B.”<sup>15</sup> In other words, Raz’s view seems to be that the normal justification thesis is implied

---

<sup>13</sup> Joseph Raz, *The Morality of Freedom* (Oxford: Clarendon Books, 1986): 46.

<sup>14</sup> Raz, *The Morality of Freedom*, 58.

<sup>15</sup> Joseph Raz, “On Respect, Authority and Neutrality: A Response”, *Ethics*, 120 (2), (2010) 299. He also means the protected reason as a whole, not just the exclusionary reason.



by more general principles of instrumental reasoning and practical rationality. For example, if I have reason to X, and treating someone's command as an exclusionary reason will facilitate doing X, I therefore have reason to treat someone's command as an exclusionary reason.

#### 4 Source 2: Non-instrumental Moral Duties

A second and distinct source of exclusion is found in certain *moral duties*. For example, the duty not to keep promises doesn't simply outweigh most of my personal prudential reasons; it excludes them altogether, so that I typically shouldn't allow myself to be motivated by them when deciding whether to break a promise. More specifically, the source of exclusion that I am aiming to isolate consists of *non-instrumental* moral duties. That it would *be the only way of reminding me to keep an important promise* is often the ground of an instrumental moral duty, one which exists in order to in some way advance other, independent moral duties.<sup>16</sup> That it would *constitute keeping my promise*, in contrast, is often the ground of a non-instrumental moral duty, at least when the promise is sufficiently important. The point of the exclusion found in the latter sort of case isn't simply to help us better act on the contents of some other consideration. Instead, the exclusion is an essential quality of the moral consideration itself—it is simply the rock-bottom, constitutive sense in which it carries weight in practical reasoning.<sup>17</sup>

Not all non-instrumental exclusionary considerations amount to moral duties, even when they do provide reasons to exclude reasons. For example, if I request my friend to not include my desires

---

<sup>16</sup> Not all considerations that advance non-instrumental moral duties involve moral *duties*, of course. In fact, not every consideration that is *necessary* to advance a moral duty is itself a moral duty.

<sup>17</sup> Raz himself refers to something like this when he says that exclusionary reason can be grounded directly (rather than instrumentally) by 'motivational considerations' (Joseph Raz, *Practical Reason and Norms*, 186). Thank you to an anonymous referee at *Law and Philosophy* for the reference.

when deciding where we should eat, I have given her a reason to exclude my desires (which would ordinarily be reasons that she should draw upon in her decision). This reason does not amount to a moral duty.<sup>18</sup> Here, I think it is important to remember that an exclusionary reason is just that—a *reason* to not treat certain considerations as reasons. Reasons are not special. Each of us has an uncountable number of ordinary first-order reasons to do an uncountable number of things. The worst that can be said of someone who fails to bring a reason to bear in deliberation about how to act, on that description of the situation alone, is that they are irrational. The same might be said for someone who fails to act on an exclusionary reason. In light of this, it is not surprising that a simple protected reason does not bind one's will in the way characteristic of moral duties. What would need to be added to this sort of exclusionary reason to turn it into something with the rarer force of a non-instrumental moral duty?

One incomplete answer is given by Raz's own understanding of obligations or duties, which differs from the one I present here.<sup>19</sup> For Raz, having a moral obligation or duty consists not only in having a protected reason (a reason to do something that also excludes some other reasons), but also one that is categorical (it has force independently from the beliefs or desires of the person to whom it applies).<sup>20</sup> I agree with Raz that moral duties (whether non-instrumental or otherwise) will be categorical in addition to being protected reasons. However, I do not think that this is *sufficient* to capture what makes them moral duties. This is shown by the earlier request

---

<sup>18</sup> I am grateful to an anonymous referee for this example.

<sup>19</sup> For discussion of this view of duties, see Raz, *Practical Reason and Norms*, 76 (though Raz uses the language of 'mandatory norms' rather than duties); Stephen Perry, "Second-Order Reasons, Uncertainty, and Legal Theory", 913-994; Timothy Scanlon, *What We Owe to Each Other* (Boston, MA: Belknap, 2000): 156-7.

<sup>20</sup> Joseph Raz, "Promises and Obligations" in *Law, Morality, and Society: Essays in Honour of H. L. A. Hart*, ed. P. M. S. Hacker and Joseph Raz (Oxford: Oxford University Press, 1977): 223. I am grateful to an anonymous referee for the reference.

example. There is very plausibly a categorical, protected reason to comply with the request, but there is no moral duty to comply with the request.

What then is distinctive about the *kind* of exclusion found in a non-instrumental moral duty? At a minimum, it has greater positive force than that of an ordinary reason to exclude some consideration. One is not merely committing a rational mistake, but doing something morally wrong, and certainly liable to moral criticism, if one fails to exclude that consideration. What, then, is the character of this moral wrong? On the one hand, it might be exactly parallel to the force of a first-order moral duty, in that it includes some positive force but also excludes some competing (exclusionary) reasons, including most or all non-moral reasons. On the other hand, it might be something rather different, since what counts as a moral duty might vary depending on the order of reasons in which we are operating. I leave it open exactly which account of this force is correct; all that is needed for present purposes is that it is in some sense more demanding than an ordinary reason.

There is one further distinctive aspect of the exclusion found in moral duties that is important to mention. Moral duties are accompanied by a second-order reason to be motivated by the duty itself rather than other considerations that speak for or against the action.<sup>21</sup> For example, if the duty of fairness requires me to contribute to a potluck, I have some reason to be motivated by the duty itself, rather than (say) matters of self-interest that happen to line up with contributing to the potluck (e.g. impressing people with my cooking skills). This does not necessarily entail that an

---

<sup>21</sup> I am grateful to an anonymous referee for suggesting that I reframe my claim as being about an *accompanying* second-order reason.

agent must always be motivated *only* by the duty, or even by the duty at all, because this second-order reason can be defeated by other considerations.

How might such non-instrumental moral duties be used to derive the existence of authoritative exclusion? We could do so by inferring the fact one person has authority over another person from some moral duty, in much the same way that we draw any inference from one moral duty to another moral duty. Roughly, we would expect a ‘principle of conservation of exclusion’ to apply when we infer a permission or requirement to exclude reasons from another permission or requirement to exclude reasons. After all, since exclusion is so rare, it seems implausible that a permission or requirement to exclude can be created out of thin air.<sup>22</sup> Take a hypothetical argument for authority based on fair play, as just one example. One rough version of the principle of fair play requires people to contribute to joint ventures from which they obtain benefits as long as they benefit from other people contributing to it.<sup>23</sup> If the joint venture in question is the law, and what people are contributing to it is their obedience, then the principle of fair play can be used to explain why people who benefit from the obedience of others to the law have a duty to obey the law themselves. Because the duty of fair play is a non-instrumental moral duty that itself excludes certain considerations (e.g. very plausibly, personal convenience should play no role in determining whether someone lives up to their duties of fairness), presumably any duty that is derived from it would likewise be a moral duty that excludes the same

---

<sup>22</sup> Note however that this seems to be going on in the normal justification thesis, which creates an exclusionary reason where there was none before. This connects to a point by David Enoch, who considers a variety of ‘conservation’ principles, including a “normativity-out-normativity-in” principle, a “reason-out-reason-in” principle, and a “duty-out-duty-in” principle. See David Enoch, “Authority and Reason-Giving”, *Philosophy and Phenomenological Research*, 89 (2) (2014): 317.

<sup>23</sup> See John Rawls, “Legal Obligation and the Duty of Fair Play” in *Collected Papers* (Harvard: Harvard University Press, 1999).

considerations.<sup>24</sup> Another source of non-instrumental moral duties that could serve as a basis for authoritative exclusion is found in personal relationships (though the relationship between these duties and moral duties is controversial). Familiarly, personal relationships are also grounds of exclusion: you shouldn't think of one's own personal interest when deciding whether to watch your child's ballet rehearsal.<sup>25</sup> This may translate into a view of political authority, as how Ronald Dworkin's view of political obligation is rooted in the thought that political obedience is a moral duty imposed upon us by the relationships that we bear to one another as fellow members of a society.<sup>26</sup>

## 5 Evaluating the Candidate Sources: Strength of Exclusion

I will next describe three necessary *features* of authoritative exclusion, and check whether these features are also found in the two candidate sources of exclusion. If they are also found in a candidate source, then that source remains a candidate for explaining authoritative exclusion. If not, that difference would seem to rule out that source from explaining how claims to authoritative exclusion are valid, at least on its own.

---

<sup>24</sup> For a particularly clear example of such a view, see Richard Dagger, "Authority, Legitimacy, and the Obligation to Obey the Law", *Legal Theory*, 24, (2018).

<sup>25</sup> Adrienne M. Martin, "Personal Bonds Without Directed Rights", *Philosophy and Phenomenological Research* 102 (1), (2019): 65-86. For further discussion of this example, see N.P. Adams, "In Defense of Exclusionary Reason", 244.

<sup>26</sup> That said, Dworkin's theory is not framed in this way, and we would have to extend it from a theory of the duty to obey the law to one of authority. Dworkin is clear that any duties that come from interpersonal relationships compete with duties that come from other considerations of justice and thus it is unclear whether he intends for them to have any exclusionary force. See Ronald Dworkin, *Law's Empire*, (Portland, OR: Hart Publishing, 2012): 202-6. See also Joseph Raz, *Ethics in the Public Domain* (Oxford, UK: Oxford University Press, 1995): 41. I am grateful to an anonymous referee for the latter reference.

The first feature of authoritative exclusion is that it is a sufficiently strong reason. It is not just that someone who is subject to authority has some reason to exclude reasons or that they should feel that they are making some sort of mistake if they do not exclude those reasons. It is that they should feel normatively *bound* to exclude those reasons. This is particularly true of legal and political authority, but extends to other kinds of authority as well.

The normal justification thesis does not guarantee that authoritative directives will be strong in this way. Consider Stephen Darwall's example of B, who has very strong prudential reason to get out of bed in the morning (like not wanting to miss class), but also has some competing reasons to stay in bed (like its warmth and coziness).<sup>27</sup> Although B overall has much stronger reason to get out of bed, he finds every morning that he is just unable to get himself to act on this stronger reason. Finally, he hires A to wake him up and order him out of bed, knowing that he will be motivated to follow A's demands. This example appears to belong to the fourth set of cases in which Raz proposes the normal justification thesis is satisfied, the ones where deciding for oneself "causes anxiety, exhaustion, or involves costs in time in resources."<sup>28</sup> B would better conform to his reasons by treating A's orders as exclusionary reasons, which implies for Raz that those orders really do give B those exclusionary reasons. Thus, according to the normal justification thesis, it appears that A would actually have authority over B. Yet it seems intuitively clear that the A does not possess any legitimate authority over B.<sup>29</sup>

---

<sup>27</sup> Stephen Darwall, "Authority and Reasons: Second-Personal and Exclusionary", *Ethics*, 120 (2), (2010): 273.

<sup>28</sup> Raz, *The Morality of Freedom*, 75

<sup>29</sup> In fact, Darwall goes a step further and suggests that there are no exclusionary reasons at all in this case.

One might object that Darwall's example does not satisfy the NJT, since B already knows that he ought to get out of bed and would thus do no better at complying with his reasons by treating A's commands as exclusionary than by relying on his own assessment of the first-order reasons. However, it is possible for someone to be more successful in acting on another person's judgment rather than their own judgment even when the contents of those judgments are identical.<sup>30</sup> Moreover, the suggestion that Darwall's example does not satisfy the NJT would drastically limit the reach of the normal justification thesis. The same sort of argument would apply not only to these akratic cases, but also to epistemic cases. As Raz himself has observed, the epistemic cases may be even less plausible than akratic cases, because after A has issued his directives, it is no longer true that A knows "what [he] should do better than [B] does."<sup>31</sup> In short, because B knows as much as B once A has directed him, he no longer needs to treat his directive itself as an exclusionary reason—he has effectively stepped out of the picture and can rely on his own judgment at that point.

One might also object that the underlying reasons in this example are not categorical, and that Raz already grants that authority will only be 'binding' if the underlying reasons are categorical.<sup>32</sup> But we have already seen from the earlier example of the request to a friend not to include one's own desires in decision-making that the categoricalness of a reason, on its own,

---

<sup>30</sup> I think that this assumption is necessary to make sense of how the normal justification thesis could be a plausible theory of legal and political authority at all. So often the law concerns matters in which we already know what to do—e.g. the law against murder. If the state is to have any legitimate authority over us with respect to these matters, it must be the case that we somehow do better in treating the law rather than our own judgment on these matters as binding.

<sup>31</sup> Raz, 'On Respect, Authority, and Neutrality,' 301. For further discussion, see Evan Fox-Decent, "Fiduciary Authority and the Service Conception", in Andrew Gold and Paul Miller (eds.), *Philosophical Foundations of Fiduciary Law*, (Oxford: Oxford University Press, 2014).

<sup>32</sup> Raz, *The Morality of Freedom*, 64. This point is raised by Ezequiel Horacio Monti, "Darwall on the Normal Justification Thesis", *Ethics*, 128, (2018): 441.

does not guarantee that it will be strong or binding. (The request applies categorically, independently of the beliefs or desires of the person to whom it is addressed, but it does not seem to give them a very strong reason to comply with the request.) Moreover, in many of Raz's own examples of the normal justification thesis, like pharmaceutical or river navigation, it is unclear whether the dependent reasons are categorical reasons, so if he were to insist that categoricity is necessary for bindingness, the normal justification thesis would generate legitimate authority far less often than he intends.<sup>33</sup>

To complete the argument of this section, note that the exclusion found in non-instrumental moral duties is (in contrast) quite strong. When one has a moral duty to keep a promise, for example, one has a strong, binding reason to exclude the contrary reasons (e.g. of self-interest) that one is supposed to exclude.

## **6 Evaluating the Candidate Sources: Range**

A second, closely related feature of authoritative exclusion is that it excludes a wide *range* of reasons, including both moral and non-moral reasons. At the farthest extreme, if there is such a thing as divine authority, it presumably excludes every reason aside from those supplied by its own directives. And the claims of political and legal authorities are almost as demanding. While it is likely that there are some kinds of reasons that they cannot exclude, it seems that these authorities claim to exclude many of the non-moral and moral reasons of their subjects.<sup>34</sup> States frequently claim to be able to regulate matters of moral controversy, like abortion and

---

<sup>33</sup> Raz, *The Morality of Freedom*, 75.

<sup>34</sup> See Susanne Sreedhar, *Hobbes on Resistance: Defying the Leviathan*, (Cambridge: Cambridge University Press, 2010).



euthanasia, and even ordinary coordination problems often raise questions of morality and justice. One might object that the law (particularly in liberal states) frequently contains exemptions for personal conscience, which one might interpret as existing to allow people to act on their moral reasons even when they conflict with the state's directives. However, unless these are somehow intrinsic to the authoritative claims themselves, this does not show that those claims cannot exclude moral reasons. The fact that political and legal authorities are morally permitted to decide the extent of those exceptions (e.g. to the extent that a constitution has morally optional provisions) seems to suggest that they are not intrinsic to political and legal authority.

The service conception has a compelling explanation of why an authority might exclude some moral reasons from consideration. If I can better act on my reasons (including moral duties) overall by excluding some of them in my deliberations about how I should act, it seems like I should do so. There is no reason to think that moral duties would be exempt from this.

In contrast, it is very difficult to see how non-instrumental moral duties would require or even permit the exclusion of other moral duties. For example, it is hard to see how a duty of fair play could require someone not to take into account other moral aspects of the actions they are commanded to do. Fair play has only contributed another moral duty to the pile of duties that a subject should deliberate about. That said, I grant that there are a few cases in which one moral duty may plausibly exclude other moral duties. For example, suppose that I could prevent two people from being tortured by torturing someone myself.<sup>35</sup> Even though intuitively the former

---

<sup>35</sup> I am grateful to an anonymous referee for this example.

consideration has greater strength (involving two people) than the latter consideration (involving one person), many would say that I still have a duty not to torture the person, one rooted in first-personal considerations of integrity rather than third-personal considerations. This overall duty could not be explained by the relative strengths of the duties, and instead must be explained by the latter having some exclusionary force over the former. Nonetheless, it is unlikely that one's own personal integrity is likely to be the ground for authoritative justification—if anything, we normally experience personal integrity as a limit on whether we should obey authorities.

## **7 Evaluating the Candidate Sources: Asymmetry**

The directives of political and legal authorities are asymmetrical in that they typically exclude reasons against performing them but not reasons supporting them. For example, the legal directive against theft is intended to give me reason not to be motivated by the considerations that speak in favour of theft, like the potential proceeds of the crime. But it does not seem to be intended to subtract from my reason to be motivated by the considerations that speak against theft, like jail time or the intrinsic moral wrongness of theft.<sup>36</sup> A person who is motivated by those independent considerations rather than the law does nothing wrong and does not seem to be irrational in any sense.<sup>37</sup>

The normal justification thesis correctly predicts that the directives of legitimate authorities will often exclude reasons in this asymmetrical manner. The normal justification thesis only excludes

---

<sup>36</sup> I grant that this does not extend to *all* cases of authoritative exclusion—a drill sergeant, for example, might reasonably expect that her commands exclude both pro- and against- reasons. However, it does apply to the central authoritative exclusion exerted by political and legal authorities. I would interpret cases like the drill sergeant, when employed by political and legal authorities, as being cases of delegated authority, where the basic authority of legal and political systems is used to create officers whose delegated authority possesses different properties.

<sup>37</sup> These considerations put pressure on the view that legal and political directives are exclusionary at all. One might think—perhaps the state is simply concerned with how we act rather than how we deliberate about how to act.

reasons insofar as doing so assists one in conforming with what there is most reason to do. This leads to a piecemeal and asymmetrical stance about which reasons are excluded. While having more motivating reasons *against* doing what is demanded would typically undermine someone in doing what they have reason to do, most of the time it couldn't hurt to have more motivating reasons *in favour* of doing what is demanded.<sup>38</sup>

The pattern of exclusion with respect to non-instrumental moral duties is considerably more complicated. In general, non-instrumental moral duties are accompanied by a second-order reason to be motivated by the duty itself, and not be motivated by other reasons, one that is not grounded in instrumental considerations. This second-order reason entails that there is equally reason to not be motivated by reasons that *align* with the duty as well as reasons that *contradict* it. I grant that this reason is sometimes defeated or even undercut entirely by other moral considerations. For example, if one makes a promise to one's mother to go to lunch with her, it seems like one can and perhaps should keep the promise because one will enjoy the lunch and not because of the duty.<sup>39</sup> In cases like this, some value in the activity would be lost if one performed the activity simply out of duty. However, examples like these do not cut against the claim that the duty holder has *reason* not to be motivated solely by the non-moral considerations; they just show that this reason is sometimes outweighed. Moreover, even if there is nothing wrong with being motivated by non-moral as well as moral considerations, I think there is intuitively *greater* asymmetry in the normal justification thesis than there is in moral duties, so the exclusion found in the normal justification thesis at least *better* reflects that found in moral

---

<sup>38</sup> With one exception being the case where deliberation is cognitively costly—e.g. traffic lights.

<sup>39</sup> I am grateful to an anonymous referee for this example.

duties.

## **8 Combining the Candidate Sources: Adding**

So far, I have identified two sources of exclusion and argued that each of them possesses some but not all of the features of authoritative exclusion, particularly the exclusion found in political and legal authority. The normal justification thesis captures the range and asymmetry of political and legal authoritative exclusion, but not its bindingness. The exclusion found in non-instrumental moral duties captures the bindingness of political and legal authoritative exclusion, but not its range and asymmetry. Altogether, it does not appear that authoritative exclusion can be explained solely in terms of either of these sources. This leaves four possibilities. First, perhaps it is justified by a third source of exclusion unrelated to practical rationality or non-instrumental moral duties. Second, perhaps authoritative exclusion is *sui generis* and unrelated to other kinds of exclusion at all. Third, perhaps philosophical anarchism is correct and authoritative exclusion is simply impossible, either conceptually or substantively. Finally, perhaps it is some *combination* of the two sources that is at work in authoritative exclusion.

Although I will not aim to fully rule out the first three possibilities, I wish to explore and give a very limited defense of the final possibility, which I think gives an attractive and plausible picture of political and legal authority. However, there are many possible ways of combining two sources of exclusion, and it is worth first describing the ones that do not work before advancing the one that I think does work.

The first I will consider involves simply *adding* two sources of exclusion together. In this case, two or more sources are present, each producing independent reasons to exclude certain reasons, sometimes in consort and sometimes in conflict with one another. We might imagine a view of authoritative exclusion, for example, that understands it as grounded partially in the normal justification thesis and partially through non-instrumental moral duties. In other words, on an additive view, a person is subject to authoritative exclusion if the normal justification thesis applies to them *and* if some moral duty (e.g. fair play) applies to them which implies that they have a duty to exclude certain reasons.

Does such an additive combination of the normal justification thesis and moral duty-based exclusion reflect the features of authoritative exclusion better than either of them taken individually? This additive view will contain some of the bindingness possessed by the exclusion found in moral duties, and so will bind its subjects in some of the way that we expect from authorities. The contribution from the normal justification thesis ensures that the additive combination of the two views will give the subject some reason to exclude both moral and non-moral reasons. However, in doing so it will necessarily give greater reason to exclude non-moral reasons than moral reasons, since it draws some of its exclusionary force from the moral duty, which typically does not exclude other moral reasons or duties. Moreover, the bindingness it does have will only apply to excluding *non-moral* reasons, since the bindingness comes from the duty-based exclusion, which I argued earlier does not typically exclude other moral reasons. Altogether, a simple additive combination of the two sources better corresponds to authoritative exclusion than either of them taken individually, but it still does not adequately replicate the features we find in authoritative exclusion.

## 9 Combining the Candidate Sources: Nesting

Another way of combining two sources of exclusion together is to *nest* them. This involves arranging them so that one source of exclusion is specified in the description of the other source of exclusion.

One initial sense in which the two candidate sources might be nested together is if there were a moral duty to comply with the normal justification thesis.<sup>40</sup> This would elevate the normal justification thesis to a moral duty to take authoritative directives as exclusionary reasons insofar as doing so helps one better conform to reasons in general.

The exclusion in this view would presumably derive its features from both the underlying moral duty and the kind of exclusion required by the normal justification thesis itself. To assess whether it is an adequate theory of authoritative exclusion, the next step would be to ask whether it possesses the features that are characteristic of authoritative exclusion. However, there is no need to proceed to this step because it is doubtful in the first place that there is a moral duty to subscribe to the normal justification thesis. The basic problem is that it is doubtful that there is any moral duty to be rational. This is borne out by Darwall's example which I mentioned earlier— if someone is not sufficiently motivated to get out of bed and go to class, it is easy to see how they are being irrational, but much harder to see how they are violating some moral duty.

---

<sup>40</sup> One might imagine a schematic (though rough) representation of this nested view as 'Moral duty {normal justification thesis}'. The contrasting view, discussed at the end of this section, would correspondingly be represented as 'Normal justification thesis {moral duties}'.

Another form of nesting involves *scope restriction*. We might restrict the scope of the normal justification thesis so that it only takes effect when authority helps those subject to it act on some *subset* of their reasons. We might expect such a view to take on some of the features of the normal justification thesis and some of the features of the subset of reasons to which it is restricted. I will consider a few different versions of possible scope restricted versions of the normal justification thesis (though note that the first two do not involve moral duties *per se*).

First, though Raz does not restrict the scope of the normal justification thesis to moral duties, he does restrict its scope to *categorical* reasons, since he grants that the legitimate authority granted by the normal justification thesis is only ‘binding’ if the dependent reasons are categorical.<sup>41</sup> However, the exclusion found in this view is unlikely to differ from the exclusion in the original normal justification thesis, since it contains no new source of exclusion apart from the normal justification thesis or instrumental reason itself. This implies, for the reasons mentioned earlier, that this view cannot satisfactorily explain the strength of authoritative commands. Even if an authority helps one better act on one’s categorical reasons, this does not imply that its commands have appropriately strong force.

A slightly more promising version of a nested view restricts the normal justification thesis so that it only aims to help its subjects better act on their *protected, categorical* reasons, ones that are *duties* in Raz’s sense (though not in the sense of ‘moral duty’ described in section 4). Unlike the first nested view, the exclusion in this view could possess different features than the normal justification thesis itself, since it draws on a novel source of exclusion (found in the protected

---

<sup>41</sup> Raz, *The Morality of Freedom*, 64.

reason itself). Nevertheless, it is still hard to see how this view could satisfactorily explain the strength of authoritative commands. Even if an authority helps one better act on one's protected reasons, this does not imply that its commands have appropriately strong force. As we have seen, not all protected reasons are strong enough to reflect the strength of authoritative exclusion.

This brings us to a third version of the nested view, one that restricts the scope of the normal justification thesis to *moral duties* in the more demanding sense understood in this paper.

Stephen Darwall describes such a view (without endorsing it) when he writes of the normal justification thesis, "But what if the reasons themselves concern moral obligations, which, if I am right, are things we are answerable (to one another as representative persons or members of the moral community) for doing?"<sup>42</sup> I will label the view in this way:

*Morally scope-restricted normal justification thesis:* The normal and primary way to establish that a person should be acknowledged to have authority over another person involves showing that the alleged subject is likely better to comply with moral duties which apply to him (other than the alleged authoritative directives) if he accepts the directives of the alleged authority as authoritatively binding and tries to follow them, than if he tries to follow the moral duties which apply to him directly.<sup>43</sup>

---

<sup>42</sup> Stephen Darwall, "Authority and Reasons: Exclusionary and Second Personal", 276

<sup>43</sup> This is adapted directly from Raz's original statement of the normal justification thesis. (Joseph Raz, "Authority, Law and Morality", *The Monist*, 68 (1985): 299). Note I am unclear whether it should apply to instrumental and non-instrumental moral duties alike. There may be some risk of 'double-counting' if it applied to instrumental moral duties.



As we will see shortly, a view like this has great promise for replicating the features found in authoritative exclusion. However, there are serious problems with this view in its present form. First, Darwall raises his own objection to it: authority requires that those subject to it be accountable to the authority for performance of the commands, and the morally scope-restricted NJT does not guarantee the presence of that accountability. Second, and less controversially, it would be very strange if the normal justification thesis distinguished sharply between moral duties and non-moral reasons. On such a view, the normal justification thesis somehow ‘turns on’ as soon as the underlying reasons are moral. If the grounds of this version of the normal justification thesis are found in instrumental rationality, it is hard to see why instrumental rationality would contain reason to distinguish sharply between non-moral reasons and moral duties; surely its dictates should apply equally to both kinds of reasons. It would be extremely *ad hoc*.

Third, the morally scope-restricted NJT cannot be grounded in instrumental rationality. Recall that Raz claims that the normal justification thesis follows from the following principle of instrumental, means-ends rationality:

- (1) “You have reason to do A, doing B... will facilitate doing A, therefore you have reason to do B.”<sup>44</sup>

---

<sup>44</sup> Raz, “On Respect, Authority and Neutrality”, *Ethics* 120 (2), (2010): 299.

Simple substitution does not give rise to something that supports the morally scope-restricted NJT. The principle that we would get if we simply substituted ‘a moral duty’ for ‘reason’ would be the following:

- (2) You have a moral duty to do A, doing B... will facilitate doing A, therefore you have a moral duty to do B.

This principle can sometimes be false.<sup>45</sup> For example, it will often be false whenever there is any other means of accomplishing A. It is only if doing B is the only way to fulfill the moral duty, and it does not clash with other moral duties, that one has a moral duty to A. Thus, the morally scope-restricted normal justification thesis is not supported by the instrumental principle identified by Raz.

The defender of the morally scope-restricted NJT might instead suggest that it is supported by some other principle of instrumental rationality, like one that says that agents have reason to take *necessary* means to their ends:

- (3) If you have reason to A, and doing X is the only means to do A, then you have reason to do X.

In this case, one might try to derive the morally scope-restricted NJT by substituting ‘a moral duty’ for ‘reason’ in (3):

---

<sup>45</sup> Note, however, that the weaker principle ‘If you have a duty to do A, doing B will facilitate doing A, therefore you have reason to do B.’ can be derived from (3) on rational instrumentalist grounds alone.

(4) If you have a moral duty to A, and doing X is the only means to do A, then you have a moral duty to X.

As we will see, principle (4) is plausible, however it is not an inference from (3) that makes it plausible. The problem is that even though duties are kinds of reasons, and everything that is true of any reason is also true of a duty, it is nevertheless fallacious to substitute ‘duty’ for ‘reason’ at *both* places in (3). We can see this by breaking the substitution into two steps. If we substitute ‘moral duty for A’ for ‘reason to A’ in (3), we get the following intermediate principle:

(3\*) If you have a moral duty to A, and doing X is the only means to do A, then you have reason to X.

Principle (3\*) is true and can validly be derived from (3), however we cannot validly infer (4) from (3\*). The truth of (3\*) only guarantees that there is *some* reason to X; it does not guarantee that that reason is a duty. So we cannot infer the morally scope-restricted NJT from this instrumental principle about taking the necessary means to one’s ends.

Altogether, the morally scope-restricted NJT does not appear to be grounded in the most readily apparent principles of instrumental rationality. While one might seek out other principles of instrumental rationality that support it, I think it is doubtful that any of them could provide a satisfactory explanation for why we should comply with the morally scope-restricted NJT. The basic problem is that though it is true that moral duties are reasons and thus are governed in part

by the same principles of rationality that govern reasons, the principles of rationality do not have the right sort of force to ground authoritative exclusion.

## **10 The Moralized Normal Justification Thesis (MNJT)**

We can overcome the previous objection to the scope-restricted version of the normal justification thesis by seeking some ground for it apart from instrumental rationality. I propose that we return to principle (4) in the previous section:

If you have moral duty to A, and doing X is the only means to do A, then you have moral duty to X.

The problem raised with the principle in the previous section was not with the principle itself, but rather the attempt to ground it in instrumental rationality. I will now suggest that it should instead be treated as an *instrumental moral principle*. On the one hand, it is not a rock-bottom intrinsic moral duty. If someone had no other moral duties, this moral duty would never trigger for them. This moral duty is an instrumental one that supports other moral duties. On the other hand, it is a moral principle rather than a rational principle. Its correctness or incorrectness must be determined by moral reflection, and it may well turn out to have exceptions that are themselves determined by moral considerations.<sup>46</sup> Call the corresponding theory of authority the *moralized normal justification thesis (MNJT)*:

---

<sup>46</sup> For example, if the only way of accomplishing a moral duty was to violate some greater moral duty, the principle would almost surely be false in that case.

One has a moral duty to accept the directives of an alleged authority as authoritatively binding if accepting those directives would make one more likely to comply with other moral duties than if one had acted on those moral duties directly.<sup>47</sup>

Note that this is actually a *doubly* nested version of the normal justification thesis—it features moral duties in two distinct ways. It suggests that we have an instrumental moral duty to treat directives as authoritatively binding if doing so helps us better act on our other, non-instrumental moral duties. The MNJT better reflects the features of authoritative exclusion than do any of the individual or combined theories discussed so far.

*Strength.* It is clear that the exclusion in the MNJT is sufficiently strong to replicate the strength of authoritative exclusion. Intuitively, if someone would better conform to their moral duties by treating commands as authoritative, and thus excluding certain reasons from action, the requirement to do so is binding and non-optional.

*Range.* The exclusion in the MNJT applies to a range of both moral duties and non-moral reasons. If the point of MNJT is to help someone better act on their moral duties, but acting on some smaller subset of their moral duties would prevent them from acting on their moral duties in general, it seems reasonable to exclude that smaller subset of moral duties.

---

<sup>47</sup> Joseph Raz, *The Morality of Freedom*, 53.

*Asymmetry.* The exclusion in the MNJT tends to be asymmetrical, just like the authoritative exclusion in legal and political authorities. The MNJT tends to exclude reasons that speak against its directives while permitting reasons that speak in favour of its directives. Presumably if one can make even better use of authoritative directives to better act on one's moral duties by factoring in the pro-reasons, but not the con-reasons, then one should do so, particularly when these pro and con-reasons are themselves moral duties.

I will conclude by gesturing at why the MNJT not only reflects the formal properties of political and legal exclusion, but also provides a substantively attractive picture of political and legal authority. It will be helpful first to more concretely illustrate how it might apply to political and legal authority. First, like the original NJT, the MNJT leads to a piecemeal conception of authority. The state will have authority over people to different degrees depending on the extent to which they are already able to act on their moral duties. The NJT and the MNJT are also similar to one another with respect to the kinds of cases that trigger them. Again, recall that the original NJT was thought to apply in five cases:

1. When the authority is "wiser."
2. When the authority has a "steadier will less likely to be tainted by bias, weakness or impetuosity."
3. When direct individual action is "likely to be self-defeating."
4. When deciding for oneself "causes anxiety, exhaustion, or involves costs in time or resources..."

5. When the authority is “in a better position to achieve... what the individual has reason to but is in no position to achieve” (here Raz seems to be referring to the role of authority in solving coordination problems).<sup>48</sup>

These cases are just as intelligible with respect to moral duties as they were with reasons in general. For example, in ordinary traffic direction cases, the NJT and the MNJT will likely reach the same conclusion: it is best to treat as authoritative the instructions of the police officer directing traffic. The NJT will take into account certain prudential reasons, like the value of getting to where one is going more quickly, but it will also acknowledge that these reasons will be overshadowed by the moral reasons at play in the situation, like the importance of avoiding harm to others. So the NJT and the MNJT would reach the same conclusion in this case. That said, sometimes the views will reach different conclusions, particularly when authorities serve only to help people act on their non-moral or prudential reasons. For example, in Raz’s standard example of pharmaceutical expertise, it is unclear whether one violates any moral duties by making one’s own decisions about pharmaceuticals, but reasonably clear that one can easily violate what one has most reason to do by doing so.<sup>49</sup> So in this case, they will likely reach different conclusion about whether the state has authority over someone.

The MNJT retains another attractive aspect of the NJT. It is a suitable theory of authority for a liberal society in which people reasonably disagree about moral questions. People with different moral views may apply the NJT in different ways, but because the NJT is drawn from

---

<sup>48</sup> Raz, *The Morality of Freedom*, 75

<sup>49</sup> Raz, *The Morality of Freedom*, 74. To the extent that there is a self-directed or other-directed duty not to take unnecessary risks with one’s own health, however, the NJT and MNJT will reach the same conclusion. Thank you to an anonymous referee for pressing me to clarify this example.

instrumental reason rather than controversial morality, those differing moral views do not prevent them from agreeing on the fact of the NJT itself. Though Raz would likely not see it in these terms, the NJT is thus a natural account of authority for a liberal society in which there is substantial disagreement about morality.<sup>50</sup> It can enjoy support in the Rawlsian overlapping consensus.<sup>51</sup> In contrast, it is much less clear that theories of authority that are based on substantive moral views, like fair play, can bind people with a wide range of moral views so that they will acknowledge themselves as being so bound. In this respect, the MNJT strikes a middle ground between views of authority based on the NJT and views of authority based on substantive moral views. Though the MNJT is rooted in a moral principle, that principle is instrumental, and so seems likely to enjoy wider support among different moral views than non-instrumental moral matters.

## **11 Conclusion**

In this paper, I have argued that neither the standard normal justification thesis nor arguments based on ordinary moral principles can single-handedly ground an account of why an agent should exclude their reasons in the distinctive way demanded by authorities (particularly political and legal authorities). This gives some support to my view that claims to political and legal authority are instead justified by a combination of these two views which I have called the moralized normal justification thesis (MNJT). The main idea of the MNJT, roughly, is that we should exclude moral and non-moral reasons when doing so helps us better act on our moral duties, and that we do so not in order to be rational, but on morally substantive grounds.

---

<sup>50</sup> Raz is a noted critic of this form of liberalism. The question of whether the normal justification thesis applies to someone, for Raz, is something that holds largely independently of their beliefs about the reasons that apply to them.

<sup>51</sup> John Rawls, *Political Liberalism*, (New York: Columbia University Press, 2005).