

Exclusionary Reasons, Virtuous Motivation, and Legal Authority

Andrew Jordan

On a commonplace model of practical reasoning, what one ought to do is just whatever it is that the balance of reasons favor. So, when deliberating about how to act, agents might ask themselves “what do I have strongest reason(s) to do?” and upon finding an answer to this question, they should act as those reasons prescribe. Joseph Raz challenges this model. He thinks that the commonplace model cannot capture essential normative features of authority, including the authority of the law.¹ In order to adequately account for authority, he claims, we need to posit second-order reasons, that is, reasons “to act for a reason, or to refrain from acting for a reason.”² Raz calls reasons to act for reasons “positive” second-order reasons, and he calls reasons to not act for reasons “exclusionary” reasons. And, Raz claims, an adequate account of authority requires exclusionary reasons since authoritative directives exclude acting on contrary reasons.³ That is, in the face of an authoritative directive, an agent should not act for reasons, even weighty reasons, that contravene the directive.

The goal of this paper is to argue that exclusionary reasons, as they figure in Raz’s account of authority, including the authority of law, are at odds with an attractive account of moral motivation. For this reason, if there is any role for exclusionary reasons to play in a sound account of objective practical rationality, it is much more circumscribed than Raz supposes. The argument proceeds as follows. In the first section I lay out some preliminary conceptual details, including the nature of exclusionary reasons, and how they compare to two other kinds of relations that are consistent with the commonplace model—outweighing, and disabling. In the second section, I argue that an attractive account of moral motivation precludes the possibility that moral reasons can be excluded in the way that Raz’s account requires. In the final section, I argue that we can account for many of the phenomenon that might be thought to require exclusionary reasons using only the resources of outweighing and disabling. So, exclusionary reasons are in tension with a sound account of moral motivation, and are not necessary to explain many of the commonplace phenomenon that might have motivated them.

I would like to thank Stephanie Patridge, Scott Hershovitz, Ken Himma, and an anonymous reviewer at the Canadian Journal of Law & Jurisprudence for many helpful comments on this essay.

1. Indeed, others have seemed willing to follow Raz on this point. See, e.g., Don Herzog, “The Kerr Principle, State Action, and Legal Rights” (2007) 105 Mich L Rev at 38.
2. Joseph Raz, *Practical Reason and Norms* (Princeton University Press, 1990) at 39. I note that there is something of an interpretive issue here. Raz could be construed not as arguing that there are exclusionary reasons that must figure in a sound account of practical rationality, but rather as providing an account of how some subjects treat authoritative norms. I’ve framed the issue in a way that assumes Raz thinks that there are exclusionary reasons. But, the arguments that follow could be reframed as a challenge to the second interpretation—that is, as claiming that subjects ought not to treat authoritative norms in that way.
3. Joseph Raz, *The Authority of Law: Essays on Law and Morality* (Oxford University Press, 1979) at 17-19.

Limiting ourselves to the tools of first order moral reasons—including such relations as outweighing, and disabling—allows us to preserve a more attractive account of the relationship between what there is strongest reason to do, what one is motivated to do, and that for which one is praiseworthy or blameworthy. In closing, I argue that we can capture the normativity associated with legal decision-making using only the resources of the first-order model.

1. Some Conceptual Preliminaries

To begin, let us consider two kinds of normative relationships: outweighing and disabling. We recognize that some competing reasons for action are stronger than others.⁴ For instance, the fact that Sam's brother has had a small emergency and could really use some help looking after his kids for a few hours this afternoon, outweighs—is a more weighty reason than—the pleasantness of watching a basketball game on the television. The outweighed reason, the pleasantness of watching the game, still has some normative weight—it is still a reason. But, supposing that Sam ought to do what there is strongest reason to do, he ought to help his brother rather than watch the game.

Another normative relationship is that of disabling.⁵ If some consideration, X, disables a reason, Y, then X figures as the explanation of why Y is not a reason, even if it would be one in some other context. To illustrate, consider that in normal circumstances, the fact that my action would be deceptive is a reason not to perform it. But, deceptiveness is not always a wrong-making feature of an activity. For instance, when playing poker the fact that bluffing would be deceptive is no reason not to do it. What we can say about this sort of case is that participating in the practice of playing poker *disables* what would otherwise be a reason—i.e., makes it a non-reason. And a person who plays poker honestly, out of a conviction that one ought not to deceive others, is exhibiting a kind of moral confusion.

But on Raz's view, there is a third sort of normative relationship that can hold amongst reasons; some reasons exclude other reasons. The exclusionary relationship sits in an odd middle ground between the more commonplace relationships of outweighing and disabling. As with an outweighed reason, the excluded reason is still a reason—it maintains its normative force. But, as with a disabled reason, the putative reason shouldn't motivate the agent in acting as she does—she should not act for that reason.⁶ To illustrate the exclusionary relationship Raz asks us to consider the case of Jeremy, a soldier. Jeremy is ordered to appropriate a van belonging to a local. He is encouraged not to do so by his friend, who

4. The outweighing relationship need not require that all reasons are commensurable—that is, assessable in terms of a single measure of weight. So long as they are at least comparable, we can get something like outweighing off the ground. In cases where reasons aren't even comparable, then it won't make sense to talk about weight. For more in incommensurability and incomparability, see Ruth Chang's "Introduction" in R Chang, ed, *Incommensurability, Incomparability, and Practical Reason* (Harvard University Press, 1997).

5. For an extended discussion of disabling, see Jonathan Dancy, *Ethics without Principle* (Oxford University Press, 2004) at 38-43.

6. At least she should not do so where the excluded reason favors action contrary to what is prescribed by the non-excluded reasons.

points out the weighty reasons not to take the van. But Jeremy reasons that it is not for him to decide what should be done, after all, orders are orders.⁷ On Raz's account, the order excludes the contrary reasons pointed out by the friend. It's not as though the friend is getting things wrong regarding what reasons there are—as he would be if the reasons were disabled. Indeed, Jeremy could very well correctly recognize those considerations as reasons. But nevertheless, he refuses to allow those reasons to determine his action in this case. That is what it is to treat a reason as excluded.

To sum up, the distinction between first-order reasons, exclusionary reasons, and disabling conditions could be understood as follows: First-order reasons stand in relation to actions and either favor or disfavor them. Exclusionary reasons, in contrast, stand in relation to the motives of an agent. They demand of an agent that she not act for certain first-order reasons. And disabling conditions stand in relation to first order reasons to turn off their normative force—that is, to make them non-reasons.

Before proceeding, it is worth warding off a potential confusion. In some sense, in any conflict of reasons where one reason wins out, that reason is, or at least produces, a reason not to act for the weaker reason.⁸ But this isn't what Raz has in mind. He believes that in some contexts the weight of first order reasons might favor some action, X, but that the reasons favoring X can be excluded. Hence, in those contexts an agent ought not to act for those reasons. This is, we might think, what is happening in the Jeremy example. The first order reasons pointed out by the friend might, all things considered, outweigh the reasons favoring taking the van. Nevertheless, Jeremy decides that he ought to follow the order and ignore the competing reasons.

What this means is that the existence of exclusionary reasons spells trouble for the commonplace model mentioned above. Once we allow for exclusionary reasons, what one ought to do can be something other than that which is favored by the balance of reasons. Consider the Jeremy example again. Suppose that the only reason to order Jeremy to commandeer the van is that it will allow a general to get to a meeting 30 minutes earlier. Suppose also, that nothing especially important hinges on the general getting there earlier—she will be late and this may be embarrassing, but that's about it. Let us suppose, also, that the inconvenience to the tradesman outweighs the benefit to the general, but that it is a reason that is excluded by the contrary order. Now ask, what should Jeremy do, and for what reason? If the order functions as an exclusionary reason, then there are two options. Either we say that Jeremy ought to act contrary to the order and refuse to take the van, or that he should obey the order. If we say the first, however, we get the result that Jeremy ought to refuse to take the van, but that since any reason to refuse is excluded, then Jeremy ought not to refuse for any of those reasons. That would be baffling. It would suggest that Jeremy ought to refuse, but that he ought

7. Raz, *supra* note 2 at 38.

8. For example, if my friend's severe headache is a reason to get him an aspirin that outweighs the inconvenience of going up the stairs to fetch it, then there is a sense in which I ought not to act for the reason that going up the stairs is inconvenient.

not to act for any of the reasons to refuse. This is to make the motivation for the act alien to the act, and that is something that a sound account of practical reasoning ought to avoid.⁹ So, it had better be the case that exclusionary reasons determine what one ought to do, the weight of first-order reasons, notwithstanding. If the order functions as an exclusionary reason, then Jeremy ought to commandeer the van. Hence, if there are exclusionary reasons then, contra the commonplace model, it isn't always the case that what one ought to do is just what the weight of reasons favors doing.

2. The Motivational Problem

So far I've articulated three different relationships between reasons—outweighing, disabling, and exclusion—and I've explained why the exclusionary relationship is at odds with the commonplace model of normative practical rationality on which a person ought to do what there is strongest reason to do. That isn't, yet, much of an objection to exclusionary reasons. After all, the commonplace model may well be mistaken. In this section, I argue that the existence of exclusionary reasons is hard to square with an attractive virtue-theoretic account of moral motivation, and a related account of moral praiseworthiness and blameworthiness. And this, I submit, is a reason to think that if there are exclusionary reasons, the role they play in a sound account of practical rationality is much more circumscribed than Raz supposes. Indeed, we should be suspicious of the claim that the dictates of putative authorities—including legal authorities—can ground second order reasons that exclude first order moral reasons.

a. A virtue-theoretic account of moral praiseworthiness and blameworthiness

In drawing out the sort of virtue theoretic account I have in mind, it will be helpful to look to Nomy Arpaly's account of moral praiseworthiness and blameworthiness that she develops in *Unprincipled Virtue*.¹⁰ Arpaly defends two theses concerning moral praiseworthiness. First, for an agent to be praiseworthy is for her to have done the right thing for the relevant moral reasons.¹¹ This observation should be familiar. It is what animates Kant's well-known grocer example.¹² The relevant moral reasons, in turn, should be understood as those considerations that are fitting given the circumstances. To get the basic idea, consider that there may be circumstances where worrying about whether one will dirty one's clothes is fitting—for instance, if there are puddles on a street and one is walking to an important business meeting—and there can be circumstances where it is not—for instance, at the scene of a very bad car wreck. And a person who is preoccupied with cleanliness at the scene of a car wreck is to that degree less praiseworthy

9. Whiting makes a similar observation. Daniel Whiting, "Against Second-Order Reasons" (2017) 51:2 *Nous* 398.

10. Nomy Arpaly, *Unprincipled Virtue* (Oxford University Press, 2003).

11. *Ibid* at 84. See also Julia Markovits, "Acting for the Right Reasons" (2010) 119 *Phil Rev* 201.

12. See Immanuel Kant, *Groundwork of the Metaphysics of Morals*, 4:397-98.

than he would be were he to have been focused solely on the needs of the victims of the crash.¹³ The fastidious would-be rescuer is simply less heroic than the rescuer who is unconcerned about the state of his wardrobe. And this is so even if the fastidious rescuer soldiers on despite his misgivings about whether his drycleaner can save his pants.

Second, Arpaly argues that an agent is “more praiseworthy, other things being equal, the deeper the moral concern that has led to her action.”¹⁴ Having deep concern is associated, Arpaly claims, with three things. First, a certain “diehard quality”—the more concern you have for a consideration, the more it takes to stop you from acting in accordance with it. Consider two agents, one who acts morally only when it is convenient, and the other who acts morally even in the face of substantial impediments. The latter is more morally praiseworthy than the first even when they do the same thing. Second, those who have moral concern have associated emotional dispositions. For instance, a person who cares deeply about his spouse’s well-being will take pride in her successes, experience joy when she is happy, and sadness when she suffers. Third, a praiseworthy person has a kind of awareness of those considerations that are morally relevant. Those who care about sexism are more aware of sexism when it occurs, for instance. The upshot of this sort of view is that being reason-responsive in the morally best sort of way is not just having a set of beliefs about what reasons there are, but involves an entire motivational complex, which constitutes an agent’s “read” on a moral situation.¹⁵ One isn’t merely in an intellectual state of believing something to be a reason, but rather one is moved in certain ways and aware of certain things. Those considerations that are most morally relevant occupy the mind of the virtuous agent.

b. Why exclusionary reasons are at odds with the virtue theoretic account of moral motivation

However, if one accepts that moral praiseworthiness and blameworthiness should be understood in terms of responsiveness to objective moral reasons, and that responsiveness to moral reasons involves a set of motivational dispositions described above, then it will turn out that exclusionary reasons are at odds with a sound account of moral reasoning. On the exclusionary reasons model, when a

13. My own view, argued for elsewhere, is that this shows that cleanliness becomes a non-reason in that context. See Andrew Jordan, “On Reasons, Evidence of Oughts, and Morally Fitting Motives” (2014) 42 *Philosophia* 391, especially 394-97, and Andrew Jordan, “Reasons, Holism and Virtue Theory” (2013) 63 *Phil Quarterly* 248.

14. Arpaly, *supra* note 10 at 84. This second claim is more controversial. See, e.g., Julia Markovits, “Acting for the Right Reasons” (2010) 119 *Phil Rev* 201. Markovits accepts the first claim, but rejects the second one. What is less controversial is that an agent would be better, *qua* agent, if she has deep concern. For my purposes, whether we want to say that an agent is more morally praiseworthy in acting when she has deep concern, or whether we merely want to say that she is a better agent when she has deep concern won’t affect the argument that follows. Rather, it will simply require reframing the issue. Instead of saying that exclusionary reasons create a tension in praiseworthy action, we will have to say that exclusionary reasons create a tension in an account of virtuous agency.

15. See, e.g., Rosalind Hursthouse, *On Virtue Ethics* (Oxford University Press, 1999) at 139. And, generally, John McDowell, *Mind, Value and Reality* (Harvard University Press, 1998) at 50-76.

person is confronted with an excluded moral reason, she should either not have the degree of concern for it that would typically accompany a virtuous outlook given its weight—e.g., a diehard quality, an awareness of its relevance, and the associated affective responses, or she should manifest these characteristics yet act contrary to those dispositions. Neither option is acceptable. The first option would entail that if moral reasons are excluded, then the traits that typify a virtuous disposition would be incompatible with the sorts of motivations an agent should have in the face of those excluded reasons. Such a result should be unacceptable to anyone sympathetic to the virtue theoretic account of moral motivation noted above.

The second option—that an agent faced with excluded moral reasons should manifest the kinds of dispositions characteristic of virtuous agency, but act contrary to those dispositions, however, creates a kind of cognitive dissonance that cannot plausibly figure into any account of virtuous agency.¹⁶ To see why, let's return to Jeremy. If Jeremy is virtuous and the excluded reason—the harm to the tradesman—is a compelling one, then Jeremy would feel the pull of acting contrary to his orders. But, because any competing reasons are excluded, he must muscle through, despite his better sentiments, and do what he has been ordered to do. In the normal case, when a morally decent person is confronted with what she takes to be a pressing moral reason, she acts. So, what we have in the case where there is an exclusionary reason—an order, say—and a pressing first order reason to act contrary to the order, is a conflict between the recognition of something as morally salient in the case, and the putative order or command. Exclusionary reasons put the virtuous agent in a position of recognizing something a reason, and hence being motivated by it, but simultaneously judging that it would be inappropriate to actually be so motivated.¹⁷

Raz himself is aware of the apparent motivational paradox that arises if we accept his account of exclusionary reasons, though it is not clear that he grasps its full implications. He allows that there is something a little paradoxical about “reasons for not being guided by reasons whose very nature is that they should guide.”¹⁸ But he tries to lessen this worry by arguing that reasons are mere candidate guides. That is, reasons are only potential guides; they need not actually motivate a person in acting as she does. As Raz puts it, reasons require conformity, not compliance. One conforms with a reason if one does what the reason favors. One complies if one also acts for that reason. For instance, if Samantha ought to

16. Many virtue theorists argue that virtuous agency is effortless in a certain sort of way. See, e.g., *Nicomachean Ethics* 1146 (a). See also McDowell, *supra* note 15. Whether virtuous agency is effortless, in many cases, exclusionary reasons would make virtuous agency cognitively discordant. It is that result that I claim is implausible.

17. Contrast what we can say if the reason is disabled. Here, since the reason doesn't have normative force, presumably a virtuous agent would not be motivated by it. The psychological discord is avoided. Admittedly, there is something of a literature on what the psychological makeup of a virtuous agent must be in the face of disabling conditions. See, e.g., Rebecca Stangl, “A Dilemma for Particularist Virtue Ethics” (2008) 58:233 *Philosophical Quarterly* 665-78. For my take on the issue, see Andrew Jordan, “Reasons, Holism and Virtue Theory”, *supra* note 13 especially 260-65.

18. Raz, *supra* note 2 at 183.

give money to a charity because doing so will alleviate animal suffering, she conforms if she gives, even if merely to impress her friends, but she complies if she gives out of concern for the animals. And Raz thinks that if reasons don't require compliance this will partially undermine one source of resistance to exclusionary reasons—the assumption that reasons ought to motivate an agent. But, as should be apparent, one upshot of the virtue theoretic account is that, contra Raz, at least moral reasons do indeed require compliance.

Raz cites three considerations that favor the view that reasons require only conformity. First, Raz notes that there may be multiple reasons for an act, each of which is sufficient to warrant acting in that way, and there is no fault in being motivated by only some of these.¹⁹ For instance, going to the party may be enjoyable for me, it may please my friend, and it may help raise money for a good cause. If I go to the party because it will please my friend, without even considering the other reasons, then on Raz's account, I have done nothing wrong, assuming that going to the party is what I ought to do.

Second, Raz argues that if we accept that reasons can be outweighed, there is nothing wrong with not being motivated by the outweighed reasons. As Raz observes, if a person could spend the afternoon visiting his mother, or alternately weeding the garden, and the weight of reasons favors visiting his mother, there is no sense in which he does something wrong if he visits his mother, and pays no mind to his garden.²⁰ Indeed, there is no sense in which he does something wrong if he doesn't give his garden even a passing thought. And so, Raz argues, it can't be the case that reasons require compliance, because there is no fault in failing to attend to an outweighed reason.

Third, Raz notes that in the case of acts that we ought not do—kill another, say—it is often most admirable to not act for a reason at all. That is, the person who never considers the possibility of killing is more admirable than the person who entertains the thought and refrains because killing is wrong. As Raz puts the point in discussing a motive to kill, “the moment we are morally motivated in such cases we are behaving in a less admirable way than those to whom the wrongful act simply does not occur.”²¹

None of these three arguments is compelling. The first—the fact that one does nothing wrong in being motivated by only some of the considerations that favor what one ought to do—seems to rely on the thought that the moral worth of an act should be assessed along a single axis—i.e., whether the act was right or wrong. There may be a point to thinking that there is a fact of the matter about what we ought to do all things considered. But thinking of moral worth as uni-dimensional, and assessed primarily in terms of whether one did what one ought to have done all-things-considered, is an impoverished picture of the moral landscape, and it encourages an impoverished picture of the contours of sound

19. *Ibid* at 180. Raz motivates the claim by reference to reasons for belief, “[t]here is nothing wrong with us just because our reasons for holding certain beliefs do not exhaust the reasons for that belief which are available to us.” And he suggests that reasons for belief should work like reasons for action in this respect.

20. Raz, *supra* note 2 at 180.

21. *Ibid* at 181.

practical rationality. Rather, we can assess an agent and her acts along many different dimensions. If one attends a fundraising event with no concern whatsoever for the cause, but only to please a friend, this might show a person to be a good friend, but also shows her not to be particularly generous. If one does it solely because she will have a good time, the act involves a kind of self-absorption. There are many senses in which an act can be assessed as praiseworthy or blameworthy. The fact that one can be morally criticisable in cases like this—those where the reasons over-determine what one ought to do, and one acts for some subset of those reasons—suggests that mere conformity is not enough, at least in many cases. We may not expect an agent to act with an eye towards *all* the reasons that might bear on a practical situation, but this fact doesn't license the further conclusion that we are unconcerned with the motivational makeup of an agent when she acts, so long as she does what she ought.

The second consideration—that there is nothing wrong with not being motivated by an outweighed reason—mischaracterizes the resistance on the part of the person who objects to exclusionary reasons.²² Raz characterizes the resistance to exclusionary reasons as arising out of the thought that it is in the nature of reasons that they guide, and “[s]urely there cannot be reasons for not being guided by reasons whose very nature is that they should guide.”²³ But, the resistance to exclusionary reasons arises not from the fact that there can be reasons that ought not to be motivationally effective—that is, that ought not lead to action—but from the oddness of being in a situation where one correctly judges that there are stronger reasons to prefer one course of action to another, but that one ought to not be motivated by what one correctly judges to be a stronger reason to act. That is, the resistance is not to the idea that there could be reasons that ought not guide—that will be true of any outweighed reason. Rather, it is to the idea that one ought not be guided by what is otherwise a sound exercise of practical reason—where one is getting the first order reasons right. Exclusionary reasons require that one shouldn't be guided by what, at the first-order level, is an exemplary exercise of practical rationality. And it is this result, not the mere fact that there could be some reasons that ought not to guide, that grounds the resistance.

And Raz's assertion that there is “nothing wrong in not being aware of, and not being motivated to act for, reasons which are overridden”²⁴ is far too strong. Of course, an agent confronted with competing first-order reasons will also be in a position of judging that something is a reason, but also judging that she ought to act contrary to that reason.²⁵ Those partial to virtue ethics often describe such situations as involving a remainder of some kind. And, indeed, it is often more fitting for an agent to act with regret, say, since that is what it means to give the

22. *Ibid* at 182-84.

23. *Ibid* at 183.

24. *Ibid* at 180.

25. Some might be more profligate in attributing reasons than I am. I am skeptical, for instance, of the existence of reasons that are such that if they figured in the motivations of an agent would show her to be more blameworthy in some regard. So, the range of cases where there are bona fide competing reasons will be fewer than one might have thought on a more permissive account of what reasons there are. See Jordan, “Reasons, Holism and Virtue Theory”, *supra* note 13.

reasons their due in cases where there is a bona fide conflict.²⁶ But that virtue-theoretic insight runs contrary to the assertion that there is nothing wrong with not being motivated by a reason that is outweighed. If I am forced to choose between saving my nephew and saving my dog, presumably I ought to save my nephew, but a complete lack of concern for my dog seems callous at the very least. Indeed, it is noteworthy that in Raz's example the reason that is outweighed is that his garden needs to be weeded. One might have thought that this was a paradigmatic example of a mere enticing reason—a reason that doesn't get us to an ought, even if there are no competing concerns.²⁷ If there are enticing reasons, and the weedy quality of one's garden is one of them, then it should not be surprising that there is nothing wrong in being oblivious to this reason. Indeed, the strength of Raz's example may depend on him having selected a reason that isn't of the right kind to yield an ought.

The third observation simply doesn't support the claim that reasons only require conformity, and not compliance. Raz's supposition seems to be that compliance with the reasons not to murder would require something like the thought "I could kill now, and I won't because there are powerful reasons not to kill." But, it is unclear why this should be so. We don't consider killing, because, for the most part, there are no reasons to kill. And there is little point in attending to the reasons not to do something when there is no reason to do it. Where killing becomes a salient possibility, we expect compliance, and not mere conformity. That is, if the possibility of killing is put on the table by another, we expect revulsion at the very idea, and not mere accidental conformity. It is unclear why we can't see this entire disposition—where one doesn't consider murder at all unless it is presented as a salient possibility—as a way of acting in compliance with reasons. Killing just isn't a salient possibility for the most part, and hence doesn't figure in the psychology of most people most of the time. But, where it becomes salient, one does display a kind of moral defect if one merely conforms to the reasons not to kill.

c. Why a parallel problem doesn't arise with outweighed reasons

A defender of exclusionary reasons may protest that competing reasons, especially important moral ones, also put a virtuous agent in the position of recognizing something as a reason, and hence manifesting the affective responses associated with virtue, yet acting contrary to those dispositions. Hence, this tension can provide no basis for objecting to exclusionary reasons, unless one is also willing to object to the existence of competing reasons. I'm not so sure. There are at least two reasons for thinking that the tension between the virtuous agent's motivational dispositions and the act that she ought to perform isn't parallel as between competing reasons, and exclusionary reasons. First, competing

26. See, e.g., Hursthouse, *supra* note 15 at 46-48.

27. See, e.g., Dancy, *supra* note 5 at 24-25 and "Enticing Reasons" in Wallace, Pettit, Scheffler & Smith, eds, *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (Clarendon Press, 2006).

first-order reasons compete over our actions, but not over our motivations. And there is a perfectly natural way of accommodating the sense in which one can motivationally give both reasons their due, while acting only as the strongest reason requires. That is just the point about regret. But it is hard to understand how one can give the first order reasons their due, while acting contrary to weight of such reasons. And the mystery is deepened when we remember that second order reasons don't favor or disfavor action at all. That is, they do nothing to show a course of action in a positive or negative light.

Second, paradigmatic examples of second-order reasons play their putative normative role independently of their content. That is, the putative exclusionary force of an order, or an authoritative directive, obtains independently of what was ordered or directed. So, if there are exclusionary reasons, there is something that must figure in the motivation of an agent so as to overcome first order reasons, but, since the exclusionary reasons are content neutral, the exclusionary reasons, *qua* exclusionary, do nothing to favor contrary action. Being content neutral, they cannot be said to characterize anything that could figure into a conception of the good. In contrast, competing first order reasons do figure into some conception of the good. So, if we accept the virtue theoretic account of moral motivation according to which the possession of virtue is a matter of having a certain motivational complex which gives reasons their due, we can account for any tension arising out of competing reasons in terms of how those considerations figure into a conception of the good. We can tell a story about how competing goods are arranged in different contexts, and what an appropriate response might be in the face of competing considerations—when to feel regret, say. Not so with exclusionary reasons, as there is no content to them, *qua* exclusionary, that might ground such a story.²⁸

d. Why we don't need exclusionary reasons to account for the range of commonplace judgements about agents who defy orders or act contrary to the law

I've argued that a conception of virtue as involving responsiveness to reasons is at odds with the role Raz attributes to exclusionary reasons. And I've argued, contra Raz, that reasons are not merely candidate guides; salient moral reasons have a grip on the mind of a virtuous person and motivate action. All pressing moral reasons, including excluded ones, would motivate a virtuous agent in accordance with their weight. But if an agent ought not to act on excluded reasons, then a virtuous agent would be motivated to do something she ought not to do. This is at odds with the virtue theoretic idea that virtue sets the standard for right action.

28. Exclusionary reasons could be coupled with first order reasons. For instance, an order may be both an exclusionary reason not to act on reasons contrary to the order and a reason to do what was ordered. *Qua* first order reason the order might figure into a conception of the good, because a duty of obedience could be part of such a conception. But in that capacity the order merely competes with other goods.

In closing out this section, it is worth considering one final line of argument Raz offers for exclusionary reasons. With regard to a person who acts contrary to the weight of first order reasons because she takes herself to be confronted with a valid exclusionary reason, Raz claims that we are conflicted.²⁹ And he suggests that the existence of exclusionary reasons best explains this sense of conflict. As Raz puts it, we experience a “peculiar feeling of unease, which will show itself when we wish to censure a person who acted on the balance of reasons for disregarding the exclusionary reason and when we have to justify someone’s acting on an exclusionary reason against claims that the person should have acted on the balance of reasons.”³⁰ And, further, he claims that as compared to a person who does wrong because she acts on a reason that is outweighed, we are more ambivalent about our condemnation of a person who acts on excluded reasons.³¹ This unease suggests that exclusionary reasons play an explanatory role in capturing certain features of common judgments about praiseworthiness and blameworthiness.

In this vein, let us consider the moral worth of a person’s act when, because of a putative exclusionary reason, she acts contrary to the weight of first-order reasons. I submit that in such cases, our responses are not univocal. That is, we don’t always feel a sense of conflict or unease about such cases. Indeed, at least some of the time, the person who disobeys an order or violates the law is more praiseworthy for doing so precisely because her having done so comes at great potential cost. Such actions are often held up as paradigmatic of courage. These persons exhibit a deeper concern for morality, because the potential personal repercussions for disobedience are so great. But it should be admitted that it isn’t always the case that one who violates an order or the law is more praiseworthy. Whether the person is more praiseworthy will depend on the facts of the case. In some cases we might say that a person who violates an order or a law exhibits a kind of hubris, even when she does so for moral reasons. But, by the same token, we might also say of a person who simply follows the order in the face of competing moral concerns that he exhibits a kind of lack of judgment, or moral blindness, or something similar. So, the question is how to account for this diverse range of judgments regarding the moral praiseworthiness or blameworthiness of an agent, and whether any such account requires exclusionary reasons. I submit that in we can capture the relevant intuitions in such cases without positing exclusionary reasons.³²

One possible way to capture these intuitions is to understand putative exclusionary reasons—orders or authoritative directives—in terms of what Don Regan calls indicator rules.³³ Roughly an indicator rule is a rule that indicates

29. Raz, *supra* note 2 at 41.

30. *Ibid* at 41.

31. *Ibid* at 45.

32. Of course, there are many possible cases, and nothing said here forecloses the possibility that the exclusionary reasons theorist could formulate a case where the positing of exclusionary reasons is the best explanation of the relevant intuitions.

33. See Don Regan, “Authority and Value: Reflections on Raz’s Morality and Freedom” (1989) 62 S Cal L Rev 995-1095.

the layout of first order reasons, but is not itself an intrinsic reason for action, or for acting or not acting for other first order reasons.³⁴ Indicator rules need not pick out first order reasons with exact precision. But, nevertheless, such rules usually should be followed, at least under normal circumstances. As Regan explains, the rule “professors should not have sexual relations with students” is a good rule to follow. It doesn’t capture anything of intrinsic moral import, but it does generally track considerations of intrinsic moral import—e.g., power differentials, and the need to be a good mentor to students, say. And we would most often do best to follow the rule rather than to try to assess the intrinsic merits of our acts themselves, since there are good reasons to believe that we won’t be especially good at doing so. Regan’s indicator rules fit well with some of the observations from the preceding paragraph. Roughly, we tend to blame those who violate the law, or who fail to follow orders precisely when we think they overestimate their ability to assess the intrinsic moral merits of their acts, but not otherwise. In cases where a person stands up to an obvious injustice, contrary law or orders don’t undermine our moral esteem for such acts. This is so because in such cases we take the agent to be justified in her decision to act contrary to the order. The idea, then, is that whether we experience unease in such cases will depend on whether a person is epistemically justified in disregarding the order or the law. Where they are, we hold them in higher regard, but where they aren’t we hold them in lower regard because their act indicates a kind of hubris. And, importantly, all of this can be captured in terms of the normativity of first-order moral reasons. Indicator rules are decision procedures that aim at getting first order reasons correct, and as Regan points out, one can revise one’s indicator rules in the face of evidence that a better decision procedure is available.³⁵ And one is free to ignore an indicator rule when it is clear that it leads to the wrong result in some case. All of this requires an agent to continue to assess and act on first order reasons. Indicator rules do not exclude such reasons. Indeed, when one adopts an indicator rule—a general rule against sexual relations with students, say—one does so precisely because of the relevant first order reasons. An indicator rule, if legitimately adopted, should represent the agent’s best effort to give first order reasons their due, in the face of recognized epistemic challenges. Consider an agent who follows an indicator rule because of limited epistemic access to the facts that constitute what she already acknowledges as a reason for action, for instance, that power differentials should be avoided in sexual relationships. In such a case, the agent isn’t excluding acting based on her assessment of first order-reasons. Indeed, it is precisely the assessment that power differentials are reason-giving that leads the agent to follow the indicator rule. Following the rule is merely a means to achieving what, at the first order level, the agent properly judges she has reason to do.

34. See generally, Don Regan, “Reasons, Authority, and the Meaning of ‘Obey’: Further Thoughts on Raz and Obedience to Law” (1990) 3 Can JL & Jur 3.

35. Regan, *supra* note 32 at 1008.

3. Why We Can Do Without Second Order Reasons

So far I've argued that exclusionary reasons are at odds with one common virtue theoretic account of moral agency. But one may still worry that exclusionary reasons have an important role to play in explaining certain commonplace normative phenomenon. In this section I argue, first, that we can capture the normative force of many putative second-order reasons claims in first-order terms. I then consider two kinds of cases where exclusionary reasons might seem appealing, and argue that in those cases we can adequately capture the relevant intuitions using only first-order resources.³⁶

a. Putative second order reasons can be explained solely in terms of first-order normative force

Raz takes it to be obvious that there can be reasons not to act for reasons. On this point, I agree. However, Raz also suggests that all reasons against acting for a reason function as exclusionary reasons.³⁷ This latter claim is false. In this subsection I argue, first, that not all reasons to not act for a reason—call them putative second order reasons—are exclusionary.³⁸ What separates exclusionary reasons from non-exclusionary reasons to not act for reasons is the trumping role that Raz ascribes to exclusionary reasons. As Raz puts it, conflicts between first order reasons and exclusionary reasons “are resolved not by the strength of the competing reasons but by a general principle of practical reasoning which determines that exclusionary reasons always prevail, when in conflict with first-order reasons.”³⁹ However, many reasons not to act for reasons are simply workaday first order reasons to be weighed against competing first-order reasons. And hence, not all reasons to not act for reasons are exclusionary in the relevant sense. Second, I argue that once we see our way to accounting for some reasons not to act for reasons in terms of first order concepts like weight, we have the tools to account for all reasons of that kind in first order terms. To put the point another way, I grant that there are reasons not to act for reasons. But I argue that where there are such reasons, we need not account for them in terms of second-order exclusionary reasons.⁴⁰ Rather, we can better understand them as just first order reasons for doing a particular kind of action—acting with a particular motive. This will require some unpacking.

36. Emran Mian discusses other kinds of cases, including coordination cases and argues that they too can be captured in first-order terms. See “The Curious Case of Exclusionary Reasons” (2002) 15 Can JL & Jur 99, especially at 108-12.

37. See Raz, *supra* note 3 at 17. And Raz, *supra* note 2 at 39.

38. Technically there could be positive reasons to act for reasons. As those aren't at issue here, for sake of simplicity I use the term “putative second order reasons” to refer to reasons to not act for reasons.

39. Raz, *supra* note 2 at 40. It is this trumping feature of exclusionary reasons that is essential to Raz's account. Without it, Raz's concept of exclusionary reasons loses any distinctive normative dimension.

40. Of course, this is consistent with there being exclusionary reasons. The point here is that if we can capture the relevant phenomena using only first order reasons we should, insofar as such an account has the theoretical virtue of being easier to square with the moral psychological account noted above.

Consider first, that putative second-order reasons are just reasons to act with certain motivations. That is, they stand in relation to an agent's motives—to what an agent takes as reason giving—not to the underlying first-order normative reason. This is what distinguishes exclusionary reasons from disabling conditions. Disabling conditions make a consideration that would otherwise be a first order normative reason not a first order normative reason. Exclusionary reasons, in contrast, leave the first order normative force of a reason untouched. They simply require that an agent not act for that reason. To see this point, consider the following: Suppose there is a hate crime statute that holds that if a person assaults another for reasons of racial animosity, that person is subject to a penalty of an additional five years in prison. It follows from Raz's view that the penalty is an exclusionary reason insofar as it is a reason not to act for a reason. But there is no first order normative reason to act on the basis of racial animosity. Hence, in this case, it would be a mistake to say that what is excluded is a first-order normative reason. What is excluded is a motivational reason—acting for reasons of racial animosity. So far, I think Raz accepts all of this; putative second-order reasons bear on an agent's motives and they leave the underlying first-order normativity untouched. So, properly understood, the putative second order reason is just a reason to do a particular kind of thing—avoid acting with a certain motive.

This brings me to the next point. Having a reason to act with a certain mental state is not different from any other reason to perform an action in a particular way, with regards the normative force of the reason provided. One can have reasons to do things thoughtfully, and reasons to avoid doing things angrily. And it is unclear why a move to mental states requires positing a distinct second order normativity. If doing something with a positive disposition will make you enjoy it more, that is a reason to do it with a positive disposition. But if doing the same thing with a negative disposition will help make your warranted contempt more evident, perhaps that is a reason to not have a positive disposition. But again, all of this can be captured in terms of first-order reasons and the related idea of weight. If that is right, then it is hard to grasp why we need to posit a distinctive kind of second-order normative status when we are dealing with acting with a particular kind of mental state, namely, acting for a reason—that is, with a certain motivation. So, at the very least it is a mistake to assume that reasons for acting for a reason are always exclusionary.

So far I've argued that at least some reasons for (or against) acting with certain motivations can be construed solely in terms of first-order normative force. Nothing about the possibility of reasons for having certain motivations requires two kinds of normativity—the normativity associated with weight, and the normativity associated with exclusion. To drive the point home, consider the following example: John is told by God that if he will perform act X out of love for his wife, God will eliminate childhood cancer, and if he does it for any other reason, She won't. Surely John has a powerful reason to perform X out of love for his wife, and not for any other reason. But, what is that reason? It's just the first order reason that doing so will cure childhood cancer. And, importantly, whether John ought to act out of love for his wife will depend on how that reason relates to

other competing reasons. If God* will destroy the heavens and the Earth if John does act X, then there are more powerful reasons not to X, and hence not to X out of love for his wife. But this is all captured in terms of the first-order weighing relationship. John has stronger reasons to not X at all, than to X out of love for his wife. The fact that John has a reason to not act with a certain motivation excludes nothing. In deciding what to do, John must compare the reason to act with a certain mental state against all reasons to do anything else.⁴¹

So, there is at least some reason to suspect that if there are reasons to act for reasons, they just have first-order normative force. The question I want to pursue in what remains is whether we can capture everything we want solely in terms of first-order normative force. In closing, I consider two possible grounds for thinking that we need the exclusionary relationship, and argue that we can capture the relevant intuitions with first-order normative concepts. Obviously, there may be other grounds for saying that exclusionary reasons are necessary to capture the relevant normative phenomenon. But I hope at the very least to have provided a model for assessing putative examples where exclusionary reasons might be thought to figure in an adequate explanation of the normative landscape.

b. Can First-Order Reasons Accommodate the Phenomenology?

One might insist that exclusionary reasons capture something important that can't be captured in terms of first order reasons and the outweighing and disabling relationships—namely, that there are certain social roles where we want to say that a person making a decision should not act for certain reasons. For instance, a judge rendering a decision should not favor one party over another because of a familial relationship with that party, or because he has a friendship with that person. And, the thought goes, this can't be captured in terms of outweighing, because the competing reason may be more powerful—e.g., familial duty might be more important than the outcome of a relatively minor contract dispute, say. So, a person who rejects exclusionary reasons is forced to say that these reasons are disabled if they want to preserve the idea that judges ought not to act for those sorts of reasons. That is, they will have to say that while familial duty might be a reason for many things, it just isn't a reason for a judge to favor one party over another in a lawsuit. But, the objection continues, this result is unsatisfying because exclusionary reasons can explain why it is intelligible to us when a person feels the pull of a consideration that runs counter to what the exclusionary reason would require (or, indeed, acts contrary to what the exclusionary reason would require). Namely, there still is a first order moral reason to favor one's daughter, say. And, the thought goes, this is something

41. Daniel Whiting argues that it is not possible to act for a reason for a reason. Whiting, *supra* note 9 at 9–11. E.g., it is not possible to act out of love for one's wife, for the sake of preventing childhood cancer. For this reason, he thinks we should reject second order reasons. I think Whiting is likely correct about this. And this would provide an additional reason to reject second-order reasons. However, for the purposes of this paper, I assume that one can act for a reason, for a reason.

that can't be captured if the reason is disabled.⁴² After all, a disabled reason is a non-reason.

However, we should not confuse acting intelligibly and acting for a normative reason. One may act intelligibly even when one fails to act for a normative reason. For an act to be intelligible, we just have to see it as responding to something that is normatively salient. It would be unintelligible, say, if someone were to give money to Oxfam because the sky was blue. But often a person responds to something that is morally salient, and thus makes it the case that the act is intelligible, even if the salient consideration is not a reason. To illustrate, consider a judge who (unbeknownst to others) is deciding a case where his daughter is the defendant. She is clearly liable, let's suppose. However, the judge grants summary judgment in her favor out of a feeling of familial loyalty. My suspicion is that most will think that the act in this case is even worse in virtue of being a kind of nepotism. Nepotism in judging is a wrong-making feature, not a right-making feature. But, nevertheless, the father's behavior is intelligible to us. We get it, but not because we think that the status of the plaintiff as his daughter is a reason to render a verdict in her favor. Indeed, we think it counts strongly against rendering a verdict at all. We get it because we think the motivation to want to aid one's daughter is very often commendable (and, frankly, to be expected), it just isn't commendable here: It isn't a reason to decide for his daughter. What does this example show? It shows an act can be normatively intelligible even if it is motivated by a reason that is normatively disabled. Remember, the fact that the defendant is the judge's daughter counts against his granting summary judgment for her. It is a wrong-making feature of his act, and hence not a reason to do it. The point, here, is that to render an action intelligible, all that is required is that the act be related to something that has a kind of general moral salience. It need not be the case that the salience takes the form of being a reason in favor of the act performed. Indeed, disabled reasons will often have this character. We get their general normative relevance. That is why it is worth pointing out that they have been disabled; they are candidates for our consideration. We would not say, for instance, that the blueness of the sky had been disabled as a reason. It just isn't even a candidate reason (for giving money to OxFam, anyway). A person who never bluffed when playing poker out of a sense of honesty is misguided, to be sure. She treats as a reason something that isn't a reason in the context. But her behavior is intelligible to us. We can imagine, indeed, her having a preference for avoiding games that involve misleading others. But this is intelligible because of the general normative salience of honesty, not because honesty is a reason when deciding whether to bluff.⁴³ If this sort of analysis iterates to other

42. Scott Hershovitz has pressed this objection in private correspondence.

43. Christian Piller makes a similar point, in service of an argument against a certain conception of moral particularism. Christian Piller, "Particularism and the Structure of Reasons" (2006) 21(2) *Acta Analytica* 87. Roughly he argues that particularists are right about reasons, but that there is a weaker notion of normative connectedness which is principled. Related points about the possibility of general conditions of normative salience which aren't reasons have been made by Pekka Vayrynen, "Moral Generalism: Enjoy in Moderation" (2006) 116(4) *Ethics* 707, and by me, see Andrew Jordan, "Reasons, Holism and Virtue Theory", *supra* note 13.

cases where we might be tempted to say that a person—because of their social role—should not act for certain reasons, there is a way to capture the relevant phenomenon without having to posit exclusionary reasons.

c. *Who Decides?*

Next, an exclusionary reasons theorist might point to the common thought that the questions “what should I do” and “who decides what I should do” can come apart. And they might contend that we need exclusionary reasons to capture some aspects of the “who decides” question. For instance, a judge may say of a statute that it is morally awful, but unfortunately it is the law, and proceed to apply the morally awful statute. And, the thought goes, there is a good explanation of this—namely that it is for the legislature, and not for the judge, to decide what the rules governing our society should be. In the face of examples like this the exclusionary reasons theorist might say that we need the category of exclusionary reasons to explain this phenomenon. After all, we need to capture both the sense in which the statute is judged to be morally awful—that’s the first order assessment—and the claim that it is not for the judge to decide—that’s the exclusionary bit. But, this too can be captured in terms of first order reasons. The question of who decides arises not just in the context of law, but also in other contexts as well. Indeed, anyone who is concerned to respect the autonomy of others, in some sense, thinks that there are decisions that should be left to another person. This is what animates our concern about paternalistic action. We recognize that people will sometimes make bad decisions, but we think that those decisions are for them to make. And the reason why is that we respect their autonomy. However, all of this can be captured in first order terms. When a friend makes a bad choice we may refuse to stop them in many cases, but in many others we may intervene. While there may be much disagreement about when, whether, or how to intervene, I submit that when we do it is a very familiar first order moral decision. It is not as though we have a well-defined zone of reasons for intervening that are excluded, but that outside that zone we are free to intervene.⁴⁴ Rather, we ask whether respect for the other person’s autonomy is overcome by the harm she is doing to herself. And it is important to note that this isn’t an all or nothing decision. We may, in the face of a friend who is starting to make bad decisions voice our concerns. And we may refuse to help him financially, and at the limit, if things are bad enough, we might intervene more forcefully, or even have the person committed. But all of this makes sense in light of the first order balancing of two competing concerns—the individuals’ autonomy on the one hand, and their well-being on the other.

Similarly, a court’s behavior can also be best assessed as balancing a number of competing concerns. I will focus on just three. There is a concern for a kind of social stability and predictability characterized by the rule of law. There is

44. Raz often talks of exclusionary reasons as though they operate only within a certain scope, within which they successfully exclude, and outside of which they do not. See Raz, *supra* note 3 at 22.

a concern, at least in democracies, for respecting the autonomy of the people expressed through their legislatures. But, there is also a concern for substantive justice. The first two concerns can look like exclusionary reasons. A judge might say in the face of a prior precedent that it is not for her to decide, and in the face of a statute that it is not for her to contradict the will of the legislature. But what this comes to is just the judgment that these sorts of things matter, morally speaking—indeed, they may matter a great deal. However, in the face of a law that is morally noxious, the Fugitive Slave Act, say, we often find reasoning that looks very similar to the paternalism case mentioned earlier—or at least we should hope to so find. At some point, a judge may narrow the reading of the statute or the precedent in order to try to give justice more of its due. But, if it is deemed noxious enough a judge might simply decide to ignore it, or find a way to deem it unconstitutional. And this is all just perfectly ordinary first order moral reasoning. As with paternalism, there is not a well-defined zone of reasons that are excluded within which the court is constrained. Rather, there are just reasons for or against certain decisions, some of which might outweigh the others.⁴⁵

Of course there are cases and there are cases, and it may be that there are some cases where we do need exclusionary reasons to explain the relevant normative phenomenon. I suspect not, however, and in this section I have hoped to show how cases that might invite an analysis in terms of exclusionary reasons can be captured in terms of first-order normativity. Indeed, given the puzzling features of exclusionary reasons I submit that before assuming that there are exclusionary reasons, we would do well to first ask whether we would be missing anything by giving an analysis in terms of first order reasons. If we can, then we should prefer that account.

4. Conclusion

I have argued that exclusionary reasons are in tension with an attractive account of the moral psychology of virtuous agency. This, I submit, should lead one to be skeptical that such reasons can figure into an account of sound moral deliberation in the way that Raz seems to suggest they do. In light of this, I have argued that we should seriously consider whether we can say what we want to say about sound moral thought and judgment using only the resources of first-order normativity. I hope to have provided at least some reason to think that we can.

45. This kind of picture is especially attractive if one adopts the kind of view on which there is no distinctive kind of legal normativity. Rather, there is just the question “what is the legal practice” and the further question “what should we do given that there is this legal practice.” See, e.g., Hershovitz, “The End of Jurisprudence” (2015) 124(4) *Yale LJ* 1160.