

Remagen, 11. Juli 2025

Liebe Teilnehmer/innen am Heidelberger Kolloquium,

das folgende ist ein Artikel zum Thema ‚Autorität‘, an dem ich zurzeit arbeite.

Mit 28 Seiten ist er, fürchte ich, zu lang, um Sie dazu zu verdammen, ihn in seiner Gänze zu lesen und mit mir zu diskutieren. Ich schlage also vor, dass Sie sich auf die ersten zwei Drittel (bis Ende Abschnitt 11 auf Seite 20) konzentrieren. Darin geht es um die ‚conceptual question‘, was Autorität eigentlich ist, und um die ‚justificatory question‘, wie sie begründet werden kann.

(Im letzten Teil [Seite 21 bis 28], den ich vorschlage wegzulassen, geht es zuerst um die ‚moral question,‘ was Autorität problematisch macht, und dann darum, was von den vorherigen Argumenten für den speziellen Fall der politischen Autorität folgt. Falls jemand zu diesen Themen Fragen hat, beantworte ich sie auch gerne im Kolloquium.)

Ich freue mich darauf, den Artikel mit Ihnen zu diskutieren!

MfG

Daniel Viehoff

Authority and Control

Daniel Viehoff (UC Berkeley)

[Draft, July 2025. Please do not cite or circulate without prior permission.]

1.

Authority – the power to decide how another is to act – is a valuable and ubiquitous tool for organizing our common life. Authority is also widely thought to be problematic, and in need of special justification. An account of authority should illuminate both features. Yet many philosophical discussions of authority prioritize the latter over the former, with distorting effects on our understanding of authority and the social practices to which it is central.

Three questions are essential to philosophical discussions of authority. The first is *conceptual*: What exactly is authority? What distinguishes it from other phenomena in its vicinity? The second is *justificatory*: What does it take for someone to have authority so understood? Because (as reflection on the first question shows) authority is a distinctly normative phenomenon, this is closely connected to questions of justification: How can one person's authority over another be justified? The third is *moral*: How can authority so understood be compatible with the moral standing of the person subject to authority? The answers to these three questions are evidently related. What authority is determines what it takes to have it, and what moral problems it poses. And what moral problems authority poses may constrain when someone actually has it. Nonetheless, these questions are distinct in their concerns, and different considerations bear on them.

This essay offers an answer to the conceptual question that, compared to many existing accounts, is rather deflationary. Authority, it proposes, is a practical tool for giving one person control over another person's actions. A's directives control how B ought to act. This account of what authority is abandons several assumptions philosophers often make about authority. These include that authoritative directives change the subject's reasons in some meta-normatively *direct* fashion; that authoritative directives impose duties of *obedience*; and that they impose *moral duties*.

This answer to the conceptual question in turn yields an answer to the justificatory question, according to which the normal way to justify A's authority over B is by establishing that there is value in A's having control over B's actions, and B has undefeated reason to contribute to the realization of that value. I show that this answer, with its focus on A's controlling B's actions rather than A's controlling B's reasons for action, escapes several problems faced by Joseph Raz's influential account of practical authority.

Finally, the focus on control helps clarify the moral problems posed by authority relations, and their bearing on the justification of authority.

I conclude by briefly sketching the implications that this account has for the type of authority with which political philosophers are especially concerned: the authority claimed, and possessed, by political and legal institutions.

2.

An account of what authority is has both a stipulative and an interpretative dimension. It stipulates which among a variety of different social phenomena it seeks to offer an account of. And it seeks to offer an account that illuminates the role that this phenomenon plays in familiar social practices. (The two dimensions are not unrelated: What motivates the stipulative focus on a particular phenomenon is ordinarily the thought that this phenomenon is important for social practices that matter to us. The plausibility of this judgment is, in turn, to be assessed by the light that the account offered sheds on these practices.)

Let me start with some stipulations. Authority is what is at issue when I direct my child to go tidy her room, and she follows my directive; when a supermarket manager orders an employee to restack shelves with pasta, not with rice; or when a judge orders a prison warden to release a prisoner at once. Each of these cases involves one person's power to make binding decisions for (or, as I will interchangeably, give binding directives to) another.

This stipulative characterization of authority – as the power to make binding decisions for another – is broad, and will be sharpened as we proceed. But it suffices to already distinguish the phenomenon under consideration from various others in its vicinity. Authority (as the power to make binding decisions for another) is distinct from the power to give directives backed by coercive threats, or the moral permission to do so. And it is distinct from the power to simply physically compel another, or the moral permission to do so.¹ Authority, coercion, and compulsion frequently co-occur in our social practices (like politics, law, and parenting). Each gives one person control over another and his actions. But they are different modalities of control: authoritative control is exercised by mere say-so, without the need for threats or compulsion. This makes authority a distinctly useful means of control, but also an especially puzzling one.

Next (and still in a rather stipulative vein), we should distinguish between two different perspectives we may take on authority as one person's power to make decisions for another.

From an external (or descriptive) perspective, A has authority over B iff A's directives control what B in fact does. Call this *de facto* authority. Note that *de facto* authority is distinguished from other kinds of *de facto* power by the particular means via which A controls B's actions: A's control over B must proceed via uptake of A's directives as settling what B ought to do, rather than via A's threats or physical compulsion (or, for that matter, bribery or persuasion). This characterization is intentionally silent on the question whose uptake does the requisite work. In core cases, B himself takes A's directive to settle what B ought to do. In other cases, third parties (but not B) take A's directives to settle what B ought to do, and exert influence on B, who responds to this influence by doing what A directs.

From an internal (or normative) perspective, A has authority over B iff A's directives control what B ought to do. Call this *de jure* authority. Here too, A's control (over what B *ought* to do, not, as in

¹ As will become relevant later, authority is also distinct from other means of shaping another's behavior, such as persuasion or bribery.

de facto authority, what B actually does) must proceed via directives with the requisite content, not via alternative means for changing the normative situation, such as threats, bribes, etc.

How are *de facto* and *de jure* authority related? On a familiar view, which I accept, A's *de facto* authority over B ordinarily depends on B's *taking* A to have *de jure* authority over B. In other cases, A's *de facto* authority over B depends on a third party C's taking A to have *de jure* authority over B, plus C's putting sufficient pressure on B to act in accordance with A's directives. So *de facto* authority is *conceptually* downstream of *de jure* authority.²

Less familiar, yet central to the account developed in this essay, is the thought that A's having control over how B ought to act (that is, A's *de jure* authority over B) is normally explained or justified by the value of A's having control, by mere say-so, over how B in fact acts (that is, the value of A's *de facto* authority over B), and B's reason to help realize that value. So *de jure* authority is (normally) *justificatorily* downstream of *de facto* authority.³

Notice that these two claims are compatible. Combining them does not (as some readers may fear) lead to problematic circularity: A has *de facto* authority over B because B takes A's having *de facto* authority to be valuable, and takes himself (B) to have reason to help realize this value by acting accordingly, viz., letting A's directives control his (B's) actions.

Finally, we must distinguish whether A has (or lacks) *de jure* authority over B from the question whether A's having control over B's actions is *morally objectionable* or not. I will call authority relations that are not morally objectionable *legitimate*, and authority relations that are morally objectionable *illegitimate*. (This is a stipulative use of 'legitimacy'. There are other normative claims plausibly associated with the term. But for the remainder of this essay, talk of *legitimate* authority will refer to authority that is *not morally objectionable*.)

It is possible that A has the capacity to control, by mere say-so, how B ought to act (and so has *de jure* authority over B), and yet A's possession or exercise of this capacity is morally objectionable (and so *illegitimate*).⁴ There are at least two ways in which this point may be overlooked. One may simply fail to distinguish between the questions whether authority is justified, and whether it is morally objectionable. Or one may recognize the distinction, but think that the answer to one

² See, e.g., Joseph Raz, *The Morality of Freedom* (Oxford: Clarendon Press, 1986), pp.26-8, for this 'downstream' view. My characterization of *de facto* authority is, however, narrower than Raz's. When characterizing institutions that have *de facto* authority, Raz highlights two features they possess: "they claim a right to rule [i.e., the power to make binding decisions for another]" and "they succeed in establishing and maintaining their rule." (26). (See also pp.27-8.) This leaves it open whether the institutions' capacity to rule rests on someone's responsiveness to their claim to *de jure* authority. Since it seems important to distinguish forms of influence that run through such responsiveness from those that do not, we have good reason to adopt the narrower characterization of *de facto* authority I offer.

³ 'Normally,' because there are cases where A has *de jure* authority over B, not because there is value in A's controlling B's actions, but because there is value in another action (such as B's coordinating with others), and A has control over the conditions under which B can realize this value, even though there is no value in A's having such control. I will return to this case in Section X below.

⁴ Much less common, but also possible, are cases in which A lacks *de jure* authority, yet her exercise of *de facto* authority over B is morally unobjectionable. Such scenarios are discussed in Jonathan Parry and Daniel Viehoff, "Instrumental Authority and Its Challenges: The Case of the Laws of War," *Ethics* 129 (2019), Sec. VI.

question also determines the other, and so they cannot come apart.⁵ I think it is clear that the two questions are distinct. And (for reasons I will discuss in Section X below) I also think that the substantive claim, that answers to these questions are aligned, should be rejected. (That is not to deny that they have bearing on each other.)

3.

The basic idea behind the account of authority that this essay develops is that *de facto* authority is (often) a valuable tool for organizing social interactions (between A, B, and perhaps third parties C), because it enables A to exercise control over B's action in a distinctly (i) *costless* and (ii) *fine-grained* fashion. Why? Because all that A needs to do to control B's action is to give B a directive with the requisite content ('Do x'). Having *de facto* authority over another comes as close as is normally practically possible to having the kind of control over another's actions that we ordinarily have over our own actions, which we control simply by forming an intention and following through on it. (I say 'normally practically possible' in deference to sci-fi scenarios in which my intentions are directly fed into another's brain as input into his deliberation.⁶)

This is relevant for addressing both the conceptual and the justificatory question. When we think about how to characterize what authority is, we should ask ourselves what matters from the point of view of those who participate in the social practices to which authority is central, and which a philosophical account of authority should help illuminate. And when we think about why one person's directive should control how another ought to act, we should ask ourselves what reason someone may have to let someone else's directive control their actions.

I develop the basic idea in three steps. First, I say something general about what A's controlling B's actions involves. Second, I explain why such control (which, remember, could be achieved not just by authority, but also by coercion or compulsion, and perhaps also by persuasion and reward) can be valuable. Third, I explain why authoritative control may be particularly valuable.

4.

⁵ Joseph Raz simply talks of 'legitimate' authority. Since his 'normal justification thesis' (NJT) is best understood as an account of *de jure* authority, this may suggest that he failed to distinguish the two questions. More plausibly, Raz took their answers to be closely related. Yet his own account has the resources to distinguish between these questions, and their answers. For discussion, see Daniel Viehoff, "Service and the Moral Problem(S) of Authority," in *Engaging Raz*, ed. Andrei Marmor, Kimberley Brownlee, and David Enoch (Oxford: Oxford University Press, 2025).

⁶ Let me flag that, if someone had this sci-fi capacity, it would be sufficiently different in nature, and carry with it sufficiently different benefits and costs, that we would want to distinguish it from the phenomenon I discuss here under the label 'authority.'

When is an action (whether A's own or B's) suitably controlled by A's directive (decision, intention, etc.)? Three features are relevant for our purposes.⁷

First, the directive represents some action as 'to be achieved or undertaken,' and thereby sets a standard of success. Success here consists in "a match between behavior and the representational content" of the relevant directive.⁸

Second, control requires not just that there is a match, but also that the match satisfy certain *modal conditions*. For A's directive, 'Do x,' to control B's action, it is not enough that B in fact does x. B's doing x must also be sufficiently counterfactual robust. Imagine A directs B to do x, B does x, but B does x only because B anyway feels like doing x, and would not have done x if he had not anyway felt like doing x. Then A's directive does not control B's action. We can distinguish two dimensions of robustness here, each of which comes in degrees. The first is a matter of how varied the circumstances can be for B to do as directed. (Imagine B lets A's directive settle what he does, but only as long as he doesn't feel a strong urge against doing so. Then A has some control over B's actions, but less than she would have if B's doing as directed were not conditional on the absence of strong contrary urges.) We may think of this as capturing the strength of authority, and the corresponding *bindingness* of authoritative directives. The second is a matter of how varied the directives can be for B to do as directed. (Imagine that B lets A's directive settle what he does, but only as long as A directs B to do x or y, but not if A directs B to do z.) We may think of this as the *scope* of the authority's jurisdiction. In the following, my focus will be on the former.

Third, control requires that B's doing x with suitable modal robustness is causally explained by A's directive. So it is *because* A directs B to do x that B reliably does x. Notice here that A's directive can explain that B *reliably* does x even if it does not explain why B does x *on a particular occasion*. (Imagine A directs B to do x. B was anyway going to do x, and therefore ignores A's directive. But if B had not anyway planned to act as A directs, B would have taken into account A's directive, and been motivated by it. Then A's directive causally explains why B does x modally robustly, even though it does not causally explain why B does x on this occasion. Such cases will become relevant when we discuss what authority requires.)

5.

Let me next turn to the value of A's controlling B's actions. (Remember that at this stage, we are not distinguishing between control via authoritative directives, coercive threats, physical compulsion, etc.)

⁷ My sketch of these three features is informed by the illuminating discussion of control in Mikayla Kelley, "A Control Theory of Action," *Australasian Journal of Philosophy* (forthcoming).

⁸ Joshua Shepherd, "The Contours of Control," *Philosophical Studies* 170 (2014), p.397.

A's having control over B's actions can be valuable, because it may enable A to achieve some ends she could not otherwise achieve, or achieve such an end more reliably, or at a lower cost, than she otherwise could.

Importantly, 'valuable' here need not mean 'valuable for A.' It can also mean 'valuable for B,' or 'valuable for some third-party C' (or indeed valuable for all of them). For whom it is valuable depends normally on the ends that are more effectively achieved via A's control of B's actions. Having control over the kitchen staff, so she can tell them what dishes to prepare and how, is valuable for the head chef who aims to make a name for herself when the Michelin inspector dines at her establishment. My having control over my young child so I can stop her from running into the street is valuable to both me and my young child. The government's having control over the speed with which drivers proceed in a residential area is valuable for the pedestrians this protects from suffering injuries.

Notice that A's controlling B's actions in the pursuit of some end may be valuable for B, not only if B benefits from achieving that end, but also when B, though he does not benefit, is required to achieve that end, and A's controlling B's action makes achieving that mandatory end less costly for B. If I have a moral duty to assist my fellow citizens, and the state, by collecting taxes, ensures that the duty is discharged, that may be to the benefit of my fellow citizens who receive assistance, as well as to my own benefit, if discharging my duty this way is less costly or burdensome than it would have been otherwise. (This is separate from the claim that B benefits from fulfilling moral duties he owes to C, and so also benefits from A's controlling B's actions when such control helps ensure that B fulfills these duties to C. I do not deny that there may be a sense of benefit in which this is true. But I simply set this aside here. In Section X below, I consider the further question what reasons or duties B may have to let A control his actions given the effects that such control would have on various people's non-moralized interests.)

How does A's having control over B's actions help A pursue her ends? A may be able to stop B from doing certain things that would be bad. (I can stop my child from running into the street and getting hit by a car.) A may be able to pursue certain ends more effectively because she acquires additional manpower: having a neighbor help me paint my house, under my direction, might cut in half the time the project takes. And A may be able to pursue certain ends that she could not otherwise have pursued, because, via her control over B, she also acquires certain skills, expertise, or physical prowess that she lacks if she acts by herself. (These greater skills, expertise, or physical prowess may simply be ones that B possesses individually, and puts at A's disposal; or they may be ones that are possessed collectively by a group of agents – B1, B2, B3 – who are jointly subject to A's control.)

These mechanisms are evidently not mutually exclusive. A head-chef's authority in the kitchen plausibly involves all three, for instance: she hasn't got the time to make all dishes; she may personally lack the specific expertise to create a particular dish, but have a sous-chef at her disposal who possesses it; and she may pursue a complex overarching plan for a succession of complementary dishes that is best realized by assigning different tasks to different staff members.

Let me emphasize, finally, that the value of A's controlling whether B does x is *not* simply the value of B's doing (or not doing) x. There is distinctive value in A's controlling B's actions (or, correspondingly, to B's actions being suitably robustly responsive to A's decisions).

Such control often has significant instrumental value. It enables A to effectively pursue complex plans involving the interlocking actions of B and others (herself included), and of adjusting B's actions to changing circumstances.⁹ Consider again the chef in the kitchen, who has control over how the sous-chef prepares the fish dish. If some guests complain that the fish is under-seasoned, the head-chef can adjust the instruction she gives the sous-chef. If the head-chef revises the menu has a whole, and it would flow better if a different fish were served, or the same fish were poached rather than fried, she can instruct the sous-chef to poach rather than fry the fish. If the sous-chef simply made the same changes himself, without being subject to the head-chef's control, it would either be sheer coincidence if those changes fit with the menu as a whole, or the head-chef would have to constantly keep an eye on what the sous-chef does to adjust her own choices accordingly. Similarly, if the sous-chef simply made the changes himself, then it would either be sheer coincidence that his dish fits suitably with the contributions made by other members of the kitchen staff, or the other members of the kitchen staff would have to constantly keep an eye on what the sous-chef (and, as a knock-on effect, what others in the kitchen) does to adjust their own choices accordingly.

Such control may, on occasion, also have non-instrumental value. For instance, a procedurally fair resolution of certain disputes may require that the issue in dispute be settled by – that is, be under the control of – some impartial third party, rather than be determined by the judgment of either of the disputants. Imagine two friends who disagree about how to share a treasure they find. There is value in having a third party decide the matter, even if her decision in fact coincides with what one of the disputants would have done anyway.¹⁰

6.

In the previous section I explained why we might value A's controlling B's actions. Here, I want to explain why we might *particularly* value A's controlling B's actions *via authoritative directives* – that is, A's possessing *de facto* authority over B (rather than other forms of *de facto* power).

That we might particularly value such control is suggested by the fact that many of us often *claim* authority: parents claim authority over their children, managers over their employees, coaches over their players, political and legal institutions over their citizens. Why does it matter to us (as

⁹ This is not to deny that, on occasion, all of this is possible without A's controlling B's actions. Participants in joint activities may be so well attuned to each other that they simply 'know' how the other acts, because their perspective on what is worth doing is so closely aligned, and they respond, and know that they will respond, to the same considerations in a similar fashion. (Think of football players who have played together for a long time. The midfielder about to pass and the offensive player looking to receive the ball may both know, instinctively, where the ball should, and will, go.) When this condition is met, the instrumental considerations in favor of one person's controlling another's actions may be absent. But, it should be obvious, this condition is seldom met.

¹⁰ The centrality of impartial arbitration to social life is recognized by, among others, Locke and Hobbes: John Locke, "The Second Treatise," in *Two Treatises of Government and a Letter Concerning Toleration*, ed. Ian Shapiro (New Haven: Yale University Press, 2003 [1690]), xxx; Thomas Hobbes, *Leviathan*, ed. C. B. Macpherson (Harmondsworth, Middlesex: Penguin Books, 1986 [1651]), xxx. For further discussion, see David Dyzenhaus, "How Hobbes Met the 'Hobbesian Challenge'," *The Modern Law Review* 72, no. 3 (2009); Daniel Viehoff, "Procedure and Outcome in the Justification of Authority," *The Journal of Political Philosophy* 19, no. 2 (2011).

parent, employee, coach, etc.) to have authority, and not just the moral permission to make coercive threats or use physical compulsion?

The natural answer is, I think, that we claim authority because, if our claim is accepted by the purported subjects (or suitable third parties), and we thus have *de facto* authority over them, we (i) get to *control* the subjects' actions (ii) *simply* by giving them a directive with the requisite content.¹¹ If B accepts A's claim to authority with regard to matter m, and A directs B to do x with regard to m, then this effectively settles that B intends to do (and, if B is a well-functioning agent who follows through on his intentions, does) x with regard to m.¹²

Authoritative control is thus distinctly (i) *costless* and (ii) *fine-grained*.

It is *costless* because all that A needs to do so that B does x is to direct B to do x. A need not (for instance) track B's performance, or provide incentives for B to perform. Thus A, by using her authority over B, can pursue ends the achievement of which she (A) cannot assess, track, or enforce, as long as she can outline the end in enough detail for B to pursue it.

It is *fine-grained* because the specificity with which A can determine B's action is simply a matter of A's capacity to individuate the relevant action in her directives, and B's capacity to recognize, and act on, that individuation. Once again, it does not depend on (for instance) A's capacity to identify whether B has performed the specific action.

These virtues come into relief when we contrast authoritative control with other familiar means for controlling what others do: by physically compelling them; coercing them; bribing them; or persuading them.

Contrast this with other means of controlling another's actions. Physical compulsion requires using one's own energy to overcome the resistance of the subject; and because it bypasses the subject's own agential or rational capacities, it cannot make use of these capacities in the pursuit of one's own ends. Coercion requires tracking performance, and imposing sanctions (or at least credibly threatening to impose sanctions) for non-performance. Bribery too requires tracking, and the expenditure of resources to reward performance. And persuasion requires identifying and putting forward considerations that S finds convincing, whether or not they are rationally relevant to the question at issue. So each of these alternative means of control is suboptimal – either more costly to deploy, or less fine-grained, or both – compared to the kind of control associated with *de facto* authority, and one can control another's actions simply by giving them directives with the appropriate content.

¹¹ These are paradigmatic cases of authority and the claim to it. Saying that they are allows for the possibility of deviant cases. But these cases are best understood relative to the paradigmatic cases and their significance. (For instance: perhaps sometimes someone claims authority simply because authority is a marker of social status or standing, and not because she has any particular interest in controlling others' actions. But the reason why authority is a marker of social status or standing is that it ordinarily gives someone control over others' actions.)

¹² Cf. Raz's discussion of the basic practical inference triggered by authority: Raz, *The Morality of Freedom*, p.28.

7.

We can now articulate and defend more fully the answer to the conceptual question that this essay offers. On the control-centered view it advances,

De facto authority is the power to control, by mere say-so, what another in fact does.

De jure authority is the normative power to control, by mere say-so, what another ought to do.

Three observations are worth making here. First, if B *accepts* that A has *de jure* authority over him, and thus that A's directives control what he ought to do, and B is responsive to what he ought to do, then A has *de facto* authority over B: she can control, by mere say-so, what he does.

Second, having control over what another ought to do is a matter of being able to change what reasons he has. But not every change in reasons is sufficient to provide control. (Requests provide reasons. But because the recipient is ordinarily free to weigh those reasons along others, the effect it has on what the subject ought to do, all things considered, is not modally robust enough to give the requester anything amounting to control.) What kind of change in reasons is required will be among the questions we discuss in the following.

Third, by 'normative power' I just mean the power to change another's normative situation by mere say-so. In recent years, a number of philosophers have adopted more demanding notions of normative powers, and have suggested that authority is a normative power in that more demanding sense. This, in turn, has been used both to delineate more sharply what counts as 'authority,' and to argue that only certain kinds of justifications succeed in establishing 'authority' (rather than something else that falls short of being a normative power in their more demanding sense).

The account of *de jure* authority (and its relation to *de facto* authority and control) I have offered helps to explain why we should resist these tendencies. More specifically, it provides the basis for rejecting three assumptions about authority that philosophers often make: (i) that authoritative directives change the subject's reasons in some special meta-normatively direct way; (ii) that the reasons created by authoritative directives are reasons (or duties) to obey; and (iii) that the reasons created by authoritative directives must be moral duties. Rejecting these most immediately bears on our answer to the conceptual question. But because each of these assumptions entails particularly demanding conditions for the justification of authority, it ultimately has important implications for our answer to the justificatory question as well.

8.

Theorists of authority have pointed out that not all ways of creating binding reasons for another amount to exercises of authority. If I step in front of your car, so that you will hit me if you don't stop, I create a reason (and indeed a moral duty) for you to stop. But, it seems intuitively clear, I do not thereby exercise *authority* over you. This remains true even if I created the duty intentionally. If I step in front of your car *in order* to create a duty to stop for you, I still don't

exercise authority over you.¹³ An account of authority, as the power to make binding decisions for another, should be able to explain what separates it from other capacities to impose binding reasons on others, even intentionally.

Some philosophers have responded to such examples by positing specific meta-normative mechanisms by which authority changes another's binding reasons. Authority, they propose, is the power to impose binding reasons *directly*.¹⁴ This proposal can be usefully divided into a negative and a positive part.

Negatively, it asserts that authority brings about normative change not by changing non-normative facts that have independent normative significance. That explains why my stepping in front of your car, though it changes your reasons, doesn't count as an exercise of authority: the change in your reasons depends on a change in the non-normative facts in the wrong way.

Positively, it asserts that, if one has authority, one can bring about normative change *without* changing non-normative facts that have independent normative significance. How so? Simply by intending to create the normative change (and perhaps communicating that intention).¹⁵ (So to justify authority is to explain why someone's merely intending to create the normative change, and communicating it, does in fact bring about the normative change.)

This account explains why stepping in front of your car to change your binding reasons is not an exercise of authority over you: the change in reasons occurs via the wrong mechanism.

Still, I think it is implausible. One set of worries, which I mention but don't treat as decisive, has to do with its extensional adequacy. Though it correctly distinguishes the car case from ordinary authority, it also seems to exclude from the domain of authority certain other instances that should be included.

The negative proposal excludes from the domain of authority cases in which someone steps up as a leader, L, in an emergency that requires a coordinated response from a group of bystanders. L directs each of these bystanders how to do their part to jointly address the emergency. Because she tells B1 to get water, B2 to get medication, B3 to keep track of an injured person's bleeding, B4 to contact emergency services, etc., each of them now ought to do as she has told them to, because (against the background of her directives) that is now the easiest way to coordinate their actions in light of what they can reasonably expect others to do. Imagine that their reasons are a function of what they reasonably expect others to do; and what they reasonably expect others to do is determined by L's directives. Because L brings about normative change via changes to non-normative facts – what people expect others to do – her power is not an instance of authority, according to the *directness* theorist.¹⁶

¹³ David Estlund, *Democratic Authority: A Philosophical Framework* (Princeton: Princeton University Press, 2008).

¹⁴ David Enoch, "Authority and Reason-Giving," *Philosophy and Phenomenological Research* 89, no. 2 (2014). See also Stephen Perry, "Political Authority and Political Obligation," *Oxford Studies in Philosophy of Law* 2 (2013).

¹⁵ "Political Authority and Political Obligation."; Enoch, "Authority and Reason-Giving."

¹⁶ Cf. Estlund, *Democratic Authority: A Philosophical Framework*, xxx, Enoch, "Authority and Reason-Giving.", xxx.

The positive proposal struggles with the case of an infant princess who (we would, I think, ordinarily say) has authority over an older servant, even though she lacks the conceptual capacity to grasp what it is (and thus the capacity to intend) to impose binding reasons or duties on him. From her perspective, all she does is tell the servant what to do, in the expectation that he will now do it.

I think that the directness account goes wrong in treating both of these cases as outside the domain of authority proper. But there may be ways for the directness account to be revised to accommodate the second example. (Perhaps the intention required to bring about the direct change is less conceptually sophisticated than is often assumed.) And though we may colloquially speak of authority in the emergency case, perhaps there are good theoretical reasons for drawing narrower conceptual boundaries if a narrower account has sufficient explanatory benefits when it comes to shedding light on the more central instances of the phenomenon at issue.

But I doubt that the directness account has these explanatory benefits. When proposing that a certain feature is essential to authority, we should ask ourselves whether that feature would be of significance to those who participate in the social practices to which claims of authority are central. And the subtle meta-normative distinction, between reasons created with and those created without changing non-normative facts, does not seem to track anything that participants in these practices have reason to care about. Why should it matter to me whether the power I have over my child allows me to impose binding reasons on her with or without changing non-normative facts? Similarly, why should that distinction matter to our legal and political institutions when they give binding directives to their subjects?

This doesn't mean, however, that we wouldn't care about the difference between being able to change reasons by giving authoritative directives on the one hand, by stepping in front of the car on the other. What matters to me as the bearer of parental authority is the capacity to exercise control over my child *simply* by giving my child directives with a particular content, rather than in other, more costly and less fine-grained, ways (such as putting myself in the way of a car). But this *simplicity requirement*, which mirrors the concern for the virtues of costlessness and fineness of grain, is in principle compatible with creating binding reasons *indirectly*, by means of changing non-normative facts, as long as *all it takes* to change these non-normative facts is to give a directive with the appropriate content.

(The directness account is, we may say, subject to an authority-shmauthority objection: why should participants in relevant social practices care about the presence of authority as the directness account defines it, as long as they have shmauthority, with satisfies the simplicity requirement but not the directness condition?)

9.

In familiar cases of *de facto* authority, if A's directive, 'Do x', makes a difference to how B acts, it does so via B's rational uptake of the directive, and the effect that it has on B's intentions. In this regard, exercises of authority stand in contrast to simply pushing people around (or, in a more sci-fi vein, manipulating their brain).

In light of this, it is often assumed that, if A has *de jure* authority over B, and A orders B to do x, then B has a duty (or reason) to *obey* A's directive to do x.¹⁷ For B to count as obeying A's directive, A's directive "itself must figure in some way in the practical reasoning of" B.¹⁸ (Obedience also requires that B actually conform with A's directive. Even if the directive figures in B's deliberation, if B does not conform to A's directive, then B does not count as obeying.)

There are different ways of spelling out how the directive must so figure. (Here are some options: B must be actually motivated by the directive when conforming to it. B must at least treat the directive as enabling the motivational force of another reason. B must at least take the directive to provide a reason, even if B's conformity is actually motivated by another, independent reason.) Whichever option is adopted, there is a gap between (i) A's directive, 'Do x,' requiring B to *do* x, and (ii) A's directive, 'Do x,' requiring B to *obey A's directive to do* x. The latter, but not the former, demands that A's directive figures in B's practical reasoning in a particular way.¹⁹

One reason for doubting that the directive imposes a requirement of obedience is that such a requirement is not reflected in paradigmatic practices of authority. Most prominently, the law tells us what to do, but does not care about what motivates us when we conform to its directives.

Another reason arises from reflecting on the relation between authority and control. Let's begin by highlighting what might make the obedience requirement seem plausible even (and indeed especially) from within a concern with control, before explaining why a concern for control does not yield the requirement of obedience.

For A's directives to control B's actions, B's actions must be suitably responsive to A's directives. In particular (and following from the earlier discussion of control), they must be positively responsive in a modally robust way: if A directs B to do x, then B must *not* do not-x, even if B would have wanted to do not-x, thought that (absent A's directive) doing not-x was rationally justified (i.e., there were undefeated reasons for doing not-x) or perhaps even rationally required (i.e., there were undefeated reasons for doing not-x, and no undefeated reasons for doing x), etc. If there is a good reason for A to have such control over B, then there is a good reason for B to let A's directives control his actions.²⁰ And to let A's directives control his actions requires being responsive to A's directives. So B has reason to be responsive to A's directives.²¹

Still, it doesn't follow that A's directive, 'Do x,' creates for B a reason (or duty) to *obey*.

¹⁷ For instance: "If [a directive is] issued by someone who has [authority], then [the directive's] recipients are bound to obey. The directive is binding on them and they are duty-bound to obey it." Joseph Raz, "The Problem of Authority: Revisiting the Service Conception," *Minnesota Law Review* 90, no. 4 (2006), p.1012.

¹⁸ Victor Tadros, *To Do, to Die, to Reason Why: Individual Ethics in War* (Oxford: Oxford University Press, 2020), p.58. See also Ch. 4 *passim*.

¹⁹ For further discussion, see Johann Frick and Daniel Viehoff, "Authority without a Duty to Obey," *MIND* 132, no. 528 (2023).

²⁰ I am making the simplifying assumption that the good reason for there being such control is a reason for B. Given the dialectic here, we can set aside other cases without loss.

²¹ As discussed below, that B has reason to let A's decision control his actions does not suffice to show that A's decision controls what B ought to do. We can set this issue aside here, however. (The argument against the duty to obey goes through even if we focus on cases in which B has a moral duty to let A's decisions control his actions.)

First, for A to have control over B's actions via her directives, it is sufficient that (i) B's actions are suitably responsive to what B ought to do, and (ii) what B ought to do is suitably responsive to A's directives. So if A's directive imposes (say) a binding reason on B to do x, and B is responsive to the binding reason, then A's directive controls B's actions. Yet neither A's control, nor B's responsiveness, need figure in the *content* of the reason that A's directive has created.

Second, it is worth highlighting that B can be suitably responsive to A's directives even if A's directives do not figure in B's deliberation in any specific way. Indeed, A may have suitable control over B's actions on a particular occasion even though A's directive does not figure in B's deliberation at all on that occasion. Imagine that B is anyway motivated to do x; and that, as long as B is anyway motivated to do x, B is unable to even hear A's directive telling him to do x. Yet if B were not anyway motivated to do x, then A's directive would reach him, and he would conform to it. In this scenario, A has control over B's actions, even though A's directive does not figure in B's reasoning. More commonly perhaps, A's directives may control B's action without themselves figuring in B's deliberation: many of us are responsive to what the law requires of us, not by deliberating about the law and its demands, but by running our plans by someone who would stop us if our plans were incompatible with the law. (Anyone working in a large organization with a general counsel responsible for ensuring that the institution acts within the bounds of the law will be familiar with this phenomenon.)

The more general point is that, though A's controlling B's actions requires B's responsiveness to A's directives, and if B has reason to let A's decisions control his actions, he has reason to be responsive to her directives, the directives themselves simply provide reasons for doing as directed. The requisite responsiveness does not require that the directive actually (rather than merely counterfactually) influence B's action. And the responsiveness to A's directive need not figure in the content of the reason or duty that A's directive creates: that the reason is created by A's directive, and B robustly conforms to that reason, is sufficient to ensure the requisite responsiveness.

10.

Discussions of authority often assume that authoritative directives impose *moral duties* on the subjects.²² Sometimes talk of 'duties' may simply serve to highlight that the directives are binding.²³ We will return below to the question what such bindingness consists in. But, for reasons that will become clear, that is less than what is required for a moral duty. Because moral

²² In fact, authority often includes the power, not just to impose, but also to rescind, or empower another to impose or rescind, requirements. For the sake of simplicity, I will simply speak of 'imposing' requirements.

²³ Raz regularly talks of 'duties' when characterizing the bindingness of authoritative directives, but not of 'moral duties.' (He invokes distinctly 'moral' duties, not when speaking about the duty to do as directed, but when explaining why the moral duty to solve urgent social problems can justify authority despite the importance of acting autonomously, or making decisions for oneself rather than having them made for one "by agents, automata, or superiors": Raz, "The Problem of Authority: Revisiting the Service Conception.", p.1016.) And as far as I can see, for Raz, talk of duties *simpliciter* is just meant to capture the idea of bindingness. ("If issued by someone who has [authority], then [a directive's] recipients are bound to obey. The directive is binding on them and they are duty-bound to obey it." Ibid., p.1012.)

duties are special in various ways, they require correspondingly special justification. So the moral duty assumption both constrains what authority is (the power to impose moral duties in particular) and how it can be justified (only in ways that suffice to establish the special things that are moral duties).²⁴

Those who take moral duties to be distinctive often assume that they have special *categorical* character. Such categoricity can be spelled out in two different (though closely related) ways. First, categorical reasons make demands on us that are independent of our personal projects, goal, or interests. Second, categorical reasons make demands on us that are ‘not up to us’ in some important way. A common explanation for the categorical character of moral duties accounts for both of these at the same time: our moral duties are ultimately grounded in the claims that others have on us. The claims of others constrain the pursuit of our own personal projects, goals, and interests; and these constraints – precisely because they serve to delineate the proper boundaries between our freedom to adopt and pursue projects and goals, and the space for others to do the same – are not up to us as individuals.

The assumption that authority imposes categorical reasons in the first sense (making demands that are independent of the subject’s own personal projects, goals, or interests) should be rejected because it is incompatible with one paradigmatic instance of authority, that which parents hold over children.

My child has non-moral reason to be under my control: reasons that are grounded in her own wellbeing, rather than in the wellbeing of others. Importantly, my authority over my child is not grounded in any *particular* project or goal that she currently has but could give up or change. It is, rather, grounded in the importance of developing and protecting her ability to pursue personal projects and goals, now and (especially) in the future.²⁵ So not all authoritative directives provide categorical reasons in the first of these two senses distinguished.

The assumption that authority imposes categorical reasons in the second sense (making demands on us that are not ‘up to us,’ or *non-optional*, as I will generally say) should also be rejected, or at least so severely qualified that none of the consequences normally assumed to follow from it do indeed follow. To resist the non-optionality assumption, we cannot rely on the case of parental authority. (Even though parental authority over children is importantly grounded in considerations of the child’s own wellbeing, the authority is not optional for the child: the child cannot opt out of it.)

But I think there are other examples that bring this point into relief. Imagine I am a member of a soccer team in a non-competitive league. (So nothing much hangs on my team’s performance.) When I play, I could adopt one of two different strategies. I could play as a brilliant individualist (swooping in whenever I see a chance to score for my team) or as a cog in the machine (doing my

²⁴ See, e.g., Jonathan Quong, *Liberalism without Perfection* (Oxford: Oxford University Press, 2011), ch.4; Tadros, *To Do, to Die, to Reason Why: Individual Ethics in War*, ch.4.

²⁵ Someone may object that, though these observations are plausible when it comes to children, they aren’t plausible when it comes to responsible adults. So authority over responsible adults must involve moral duties. But notice that, even if we were to grant this claim, it no longer rests on a plausible answer to the *conceptual* question. Instead, it must be based on further normative or moral considerations about the special standing of adults, and how that moral standing is incompatible with another’s authority grounded in the adult’s wellbeing.

assigned part in the collective effort, covering the opponent player assigned to me, etc.). (This is, roughly, the distinction between what German soccer commentators call *Künstler* and *Wasserträger*.²⁶)

To be a successful cog in the machine, I must let the team's coach control what I do. (Her telling me to cover the opposing team's Nr. 10 settles that I will cover them. Her telling me to fall back and protect my own goal against a sudden counterattack when everyone else is moving forward settles that I will fall back, no matter how great a contribution I could have made to our offensive play.) I cannot be a cog in the machine without letting her directives settle what I do; and if I opt for that role, then that is what I have reason to do.

And yet the project that grounds the coach's authority – my playing as a cog in the machine rather than a brilliant individualist – is something that is up to me. (My team mates have no *moral* entitlement that I play as a cog in the machine rather than as a brilliant individualist.) So my reasons to let the coach's directives settle what I do are *optional*, and I can revise or abandon the project that grounds them. Still, as long as I haven't revised or abandoned it, I have reason to let the coach's decisions control my actions. Failing to let her directives settle what I do would not be morally wrong. But it would be constitutively incompatible with the (optional) project of playing as a cog in the machine.

Coaching relations may seem non-paradigmatic instances of authority. Yet I think what I just said about the authority of the soccer coach also applies to many authority relations in employment contexts (which, I think, cannot be so easily treated as peripheral instances of authority in a society like ours). Philosophers often mistakenly assume that, by accepting a job, the employee promises to obey their superior. But much more commonly, what the employee promises (if he promises anything) is to do a decent job. And just as in football, so in employment, doing a decent job can mean quite different things. To invoke two familiar caricatures: It can be a matter of being a diligent, obedient team-member. Or it can be a matter of being an entrepreneurial rule-breaker. Neither of these is morally required just because one has accepted a job. But one of them constitutively requires letting the superior's decision control the employee's actions at work in ways in which the other one does not.

These examples cast doubt on the assumption that authoritative directives impose moral duties. They also provide support for two positive claims.

The first has to do with the role that appeals to authority play in our practical reasoning and common culture. Most commonly, authority relations are perceived as restricting the subject's freedom: the reasonable choice how to act that the subject would otherwise have had is gone, because the authority's directive now settles what the subject ought to do. (This concern looms very large in philosophical discussions of authority in general, and political authority in particular: how can authority be compatible with the subject's status as a free person?) But authority does not only play this *restrictive role*. In addition to restricting the subject's options and limiting freedom that he would otherwise have had, it also provides justification for the subject's actions and make rationally defensible actions that would otherwise have been indefensible. (Hanging back in my own half when my presence before the other team's goal would have been

²⁶ <https://www.dfb.de/trainer/artikel/verschiedene-spielertypen-so-bilden-sie-ein-echtes-team-2728>

extremely valuable may have been a sign of bad judgment or laziness if I play as a brilliant individualist. But if I play as a cog in the machine, and following the coach's directives, I have all the justification I need for staying behind. If my teammates challenge me for staying behind, 'where were you during our attack?', I have an answer if I play as a cog in the machine that I wouldn't have if I played as a brilliant individualist.) This *justificatory role* of authority relations is often neglected. But it is important for many complex social relations, including many – such as those among public officials and citizens – central to our political life.

The second has to do with the *bindingness* essential to authority relations. I have suggested that the coach (or the manager) has authority, and can bind the player (or employee), yet that authority is optional for the player (or employee) insofar as letting the coach's (or manager's) directives control his actions is one among several different ways in which he could do his part on the pitch (or at work). So what does the bindingness of authoritative directives consist in, if not in their categoricity (in either of the two senses discussed earlier)?

We can make progress on this question by returning to the earlier discussion of control. If control is central to authority, we can ask what it takes for A to have the control over what B does that counts as having authority over B. Control requires positive responsiveness as well as a certain form of counterfactual robustness. This explains how the coach can have the capacity to *bind* the player even though her authority over him is optional (in the sense discussed earlier). As long as the player plays as a cog in the machine (rather than as a brilliant individualist), the player does what the coach directs even in the face of various countervailing considerations. So: if A directs B to play defense, B plays defense even if he would rather play offense, his playing offense is (he thinks, rightly) better for his team, etc. (The broader the range of such considerations, the more robust the control that A has over B.)

To say that A's directive is binding for B is then just to say that, if A directs B to do x, then B has reason to do x even in the face of various countervailing considerations that would otherwise have led him to do not-x. (We may say that A's decision settles what B does. More technically, for A's directive to control B's choices, A's directive is not just a reason for doing as directed, but also a second-order reason for doing x in the face of various countervailing considerations.²⁷)

The robustness required for control can come in degrees. This explains, then, why the coach can have *optional authority*: The coach's capacity to bind the player is robust relative to some considerations, but not others. It is robust, in particular, relative to the player's own preference to play offense or defense; or the player's judgment (correct or false) that his playing offense or defense would be better for his team. But it is not robust relative to the player's choice whether to play as a cog in the machine, or as a brilliant individualist.²⁸

²⁷ We can usefully think of this along the lines of Joseph Raz's notion of an exclusionary reason: a reason that excludes, rather than outweighs, other considerations. Where I disagree is with Raz's characterization of exclusionary reasons in terms of deliberation. What makes authoritative directives to do x binding is not that they require us to deliberate in a certain way about doing x, but that they settle how we ought to act by excluding certain reasons from among those that could justify doing non-x. For related discussion, see N.P. Adams, "In Defense of Exclusionary Reasons," *Philosophical Studies* 178 (2021).

²⁸ One wrinkle is worth mentioning: there is a constitutive constraint on the very possibility of playing as a cog in the machine. If the player chooses to play as a cog in the machine, but then reverts to playing as a brilliant individualist

11.

As I observed in passing in Section 8, sometimes A's power to impose a binding reason on B by giving a directive can be explained by A's power to change non-normative facts which in turn impact B's binding reasons. (That was the case of the bystanders whose coordination was required to address an emergency, and whose reasonable expectations could be determined simply via a leader's directive.) In such cases, there is no deep explanatory challenge why A's directive imposes a binding reason on B, or how A could have authority over B (assuming, as I am now in light of the previous discussion, that the power A has is indeed an instance of *authority*).

But often there is such an explanatory challenge, because A's directive does not in any obvious way change the non-normative facts on which B's normative situation depends. Still, we can explain why A has authority over B in such cases: what explains A's power to give binding directives to B (if A in fact *has* justified authority over B) is that (i) *there is value in A's being able to control B's actions by mere say-so*, (ii) *B has reason to help realize that value, by letting A's directives control what he (B) does*, and (iii) *this reason is undefeated*.

This is a very general justificatory condition, which can be satisfied in various ways. A's control over B may be valuable to different people; and its value may take different forms (instrumental, constitutive, intrinsic).

A's control can be valuable for a wide range of persons. It could be valuable for B, because B's interests are better pursued (or a duty he, B, has to another is fulfilled at a lesser cost to himself) if his actions are controlled by A. It could be valuable for a third party C, because C's interests are better advanced or protected if A controls B's actions (and thus stops B from harming C, or helps B to assist C). Or it could be valuable for A, because A can better pursue her own interests if she can deploy B's actions in their pursuit. (Indeed, A's control could be valuable because it realizes some impartial value, such as protecting the Rocky Mountains, even if it doesn't benefit anyone.) What bears emphasizing (somewhat repetitively) is that we must not assume that, because A's control over B is valuable, it is valuable for A. (Nor, for that matter, can we assume that it must be valuable for B.)

A's control can be valuable instrumentally, constitutively, or intrinsically. Instrumental cases may be the most common ones: A's having control over B's actions helps bring about some other, independently valuable ends, such as reducing the risk that B will inflict harm on someone else, or increasing the chance of making a delicious and complex dinner. But there are also familiar non-instrumental cases. For instance, A's controlling B1 and B2's actions may be valuable because it enables the peaceful and impartial resolution of a conflict between B1 and B2. (We may think of the authority of an impartial arbitrator along these lines.²⁹) Indeed, at the extreme (and only very rarely in practice), A's having control over B may be valuable as such (and sufficiently so to provide an undefeated reason for B). Imagine, for instance, that A has ALS (or,

whenever it suits him, then the coach's control over him even when he plays as a cog in the machine may lack the robustness required to make him a valuable cog in the machine.

²⁹ Viehoff, "Procedure and Outcome in the Justification of Authority."

as it is often called in the US, Lou Gehrig's Disease), and is unable to move her own body. She has a weighty autonomous interest in controlling her environment. Since she cannot do so directly, via her own body, she must do so indirectly, via control over another's actions by her mere say-so. If A's interest is weighty enough, and the burdens on B sufficiently small, then this may suffice to justify A's authority over B.

This last example highlights the importance of parts (ii) and (iii) of the justification I outlined. That there is value in A's having control over B can justify authority over B only if B has reason to do his part in realizing this value. That reflects the fact that authoritative control runs through the actions of the subject.

But – and this takes us to part (iii) – being under A's control may be costly to B, and B may have reason not to bear this cost. (These reasons could be non-instrumental, sounding in autonomy or relational values. Or they could be instrumental, because being under A's control may prevent B's pursuit of his own personal projects and goals. I will say more about some of those in Section X, when I discuss the moral problem of authority.) B's being under A's authority may also impose burdens on third parties (C), and B may have reason not to impose these burdens on them.

If these countervailing considerations are important enough to defeat the reasons in favor of letting A's decisions control B's actions, then A does not have justified authority over B.

If B has not just any old reason, but a *moral duty* to let A's directives control his actions, and that moral duty is undefeated, then B is morally required to do as A directs.

If the reasons in favor of letting A's directives control B's actions are not moral duties, and they defeat B's countervailing reasons, then B is rationally (but perhaps not morally) required to let A's directives control his actions.

And if reasons in favor of letting A's directives control B's actions are undefeated, but so are reasons in favor of doing something that is incompatible with letting A's directives control B's actions, then B is free to choose between acting on either of these undefeated reasons. So letting A's directives control his actions is optional for B (just as it is for the soccer player or the employee in the earlier discussion).

Reflecting on these different justificatory scenarios helps us clarify the varieties of authority (and the importance of moral duties for the justification of authority) discussed in Section 10.

First, we should recognize that B's having an undefeated reason to let A's directives control his actions does not *suffice* to establish A's *de jure* authority. After all, B may *also* have an undefeated reason to do something other than let A's directives control his actions. If so, it is up to B to decide whether to let A's decision control his actions, and thus whether to do what A directs. But then A's decision does *not* control what B ought to do (and so A does not have *de jure* authority over B). There is, in short, a gap between 'B has undefeated reason to let A's decision control his actions' and 'A's decisions control what B ought to do.'

One might think that this contradicts the earlier observation that there can be optional authority. If the coach has *de jure* authority over the player, then, it might be thought, the player must have done something to turn the optional reason (to play as a cog in the machine rather than as a brilliant individualist) into a non-optional one – he must, say, have promised or committed to abiding by the coach’s authority.

But this is not the only way to address the gap I flagged. If the player (i) has undefeated (but optional) reason to let the coach’s decision control his actions on the pitch, and (ii) *in fact lets the coach’s decisions control his actions on the pitch*, then the coach’s directives *also* control what the player ought to do on the pitch. They control what the player ought to do as long as the player justifiably takes for granted that he will play as a cog in the machine (or, more generally, as long as the subject takes for granted that he will pursue his ends via the particular means of letting A’s decisions control his actions, and his doing so is justified). This is compatible, however, with the player’s retaining the option of giving up this means, and adopting a different one (such as playing as a brilliant individualist). So the authority remains optional in the sense flagged earlier.

Next, consider the second type of scenario, in which the reasons in favor of letting A’s directives control B’s actions are not moral duties, and they defeat B’s countervailing reasons. Here B is rationally (but perhaps not morally) required to let A’s directives control his actions. Since B has no rational choice but to let A’s directives control his actions, A’s directives control what B ought to do. So this second scenario is in principle sufficient to establish A’s *de jure* authority over B. (In particular, unlike in the case of optional authority, A’s having *de jure* authority over B is not conditional on B’s actually letting A’s decisions control his actions.) The paradigmatic instance of such authority is that of parents over children (which is neither conditional on the child’s actually adopting the parents’ authority as a means, nor grounded in the child’s moral duties to do so).

Some philosophers resist this possibility because they think it yields implausibly extensive authority relations. Consider an example often used to convey this worry: My financial advisor, who knows better than I do how to invest money, tells me how to invest my money. I will do better, prudentially speaking, if I invest as my financial advisor tells me to. So does my financial advisor have authority over me regarding my investment choices? A common thought is that the answer is ‘no,’ *even if* letting my advisor’s decision settle how I invest is the only course of action I have undefeated reason to undertake.³⁰ But then, it is thought, the second scenario cannot suffice to justify authority. Only moral reasons or duties suffice for that purpose.³¹

I agree that the financial advisor does not (without more) have authority over me in this case. But this can be explained by observing that I am not in fact rationally required to let the advisor’s decisions control my investment choices. Even if doing so will lead to the best results when it comes to my investment returns, I am not rationally required to maximize my returns. So it remains open to me to opt for a different way of investing. (I may, for instance, decide to invest in light of values other than maximal financial return: I might want to abstain from dirtying my

³⁰ This example is due to Stephen Darwall, "Authority and Reasons: Exclusionary and Second-Personal," *Ethics* 120 (2010).

³¹ Some objections along these lines rest on implausible assumptions about what authority conceptually requires. Darwall, for instance, assumes that authority requires a claim right to rule, or some second-personal claim against the subject. As the case of parental authority shows, that is clearly mistaken, and so I set it aside here.

hands by investing in unethical enterprises; I might want to be a small shareholder in my favorite football club; etc. And I might change my mind about what these values are, or what matters to me in the investment of my money.) Consequently, the financial advisor stands in the same relation as the coach: her authority over me, if she has it at all, is optional, and depends on my not just having undefeated reason to let her directives control my investment choices, but my actually letting her directives play this role.

The difficult question this raises is whether, for adults in full possession of their capacities, *all* non-moral choice situations are like this: Their interests are *always* sufficiently dependent on their actual personal projects and goals, and remain open to revision. So arguments grounded in the subject's interests *never* establish suitably robust control over what the subject ought to do, and so fall short of establishing *de jure* authority unless more is added.

One reason for doubting that this is true is that this argument clearly doesn't apply to children, and the difference between children and adults is one of degree, not kind. (Sometimes philosophers write as if it were a difference in kind, because children lack essential rational capacities that adults possess. But this generally reflects either an undue focus on very young children, or an implausible picture of older children – say, ten- to fourteen-year-olds. So if parents still have authority over older children, as I think they do, we must ask what justifies it, and why authority over adults cannot be justified on similar grounds.) Another is that there are certain all-purpose-goods that are important for the pursuit of just about any of our personal projects and goals, and protecting our access to those might be sufficiently important to justify restrictions on our freedom to revise some of our personal projects and goals. (This was the point briefly made in Section 11, when discussing whether political authority can be grounded in the interests of those subject to it.)

It is against the background of this difficult question that we can appreciate the temptation to insist that authority *must* rest on moral duties: if competent adults cannot be rationally (but non-morally) required to let another's decisions control their actions, then authority over adults is either optional (and thus in need of further 'beefing up') or grounded in moral duties. But we can also see why it would be philosophically mistaken to make authority conditional on moral duties: because, first, there is optional authority (which, though conditional on certain further conditions, is not conditional on any moral duties); second, there is, at least for non-adults, authority grounded in reasons other than moral duties; and, third, whether the latter is also true for adults is an open question, which the assumption, that authority must rest on moral duties, prematurely forecloses.

12.

It is worth contrasting the account developed here with Joseph Raz's influential account of practical authority, with which it shares certain features, and consider how it addresses certain objections that have been raised against Raz's account.

According to Raz's 'normal justification thesis', "the normal way to establish that a person has [*de jure*] 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, rather than trying to follow the reasons why apply to him directly.”³²

One objection often raised against Raz is that it struggles to explain how A’s mistaken directives can be binding for B. (This objection tracks, not Raz’s formulation of NJT, but rather his discussion of the conditions in which NJT is satisfied. So it is possible that the answer to the object that I raise is compatible with Raz’s NJT.) Yet intuitively such mistaken directives can be binding. (Raz accepts as much.)

The (alleged) problem for Raz is that, if A’s directive, ‘Do x’, is mistaken and doing not-x would be better, then B will not better conform with the reasons that he has by abiding by A’s directive than by not abiding. To address this worry, Raz highlights that B may be in no position to (reliably) distinguish mistaken directives from correct ones. So B will generally do better conforming with all of A’s directives rather than unreliably second-guessing them.

But this answer, Raz’s critics worry, is inadequate. Grant that B will do better cultivating a general disposition or strategy to follow all of A’s directives rather than only some. Still, it isn’t the case (these critics object) that B will better conform with the reasons he has by actually doing x when A directs him to do x. That we have reason to cultivate a disposition does not show that we have reason to act in accordance with it in particular instances.³³

This objection can, however, be answered. (This answer may in principle be available to Raz as well. Still, it becomes, I think, more salient against the backdrop of the control-centered theory developed here.)

The objection rests implicitly on the assumption that the reason B has to follow A’s directive to do x is closely tied up with B’s (directive-independent) reason to do x. But this overlooks that among the reasons that B has is a reason to let A’s directives control his actions. The value of A’s controlling what B does is not coextensive with the value of B’s doing what A directs. (This mirrors the point that control requires modal robustness, and that there is value in A’s controlling B’s actions.) So there may be value in, and reason for, letting A control B’s action (including his action to do x) even if there is no value (independent of such control) in, or reason for, B’s doing x. (What is the value of A’s controlling B’s actions, over and above B’s doing what A directs? It is, for instance, the value of A’s having the capacity to pursue complex plans when she has modally robust control over the actions of others.)

This general point, about the value of control, helps clarify how Raz’s observation, about the difficulty of distinguishing good and bad cases, justifies the bindingness of A’s directive without conflating the reason for adopting a disposition or strategy and the reason for doing something in accordance with the disposition or strategy. Grant that B cannot reliably distinguish between A’s correct and erroneous directives. By following A’s directives only if he takes them to be non-erroneous, B makes A’s control over his actions less robust than it would otherwise be. This is a statement, not about B’s general disposition or strategy, but about any particular instance in

³² Raz, *The Morality of Freedom*., p.53.

³³ Quong, *Political Philosophy: The Puzzle of Legitimate Injustice*., pp.xxx.

which the question arises whether to let A's directive control his actions. So if there is value in letting A's directives robustly control his actions, and (trying to) distinguish between correct and erroneous directives undermines such robustness, then B has reason not to (try to) distinguish between correct and erroneous directives. Notice that the argument here is not that B's incapacity to distinguish good and bad directives *explains* why B has reason to let A's directives control his actions. We need independent grounds for explaining why B should let A's directives control his actions. (This could, but need not, be the fact that B does a bad job distinguishing, not A's good or bad directives, but rather cases in which he, B, has good reason to take one or another course of action.) B's incapacity to distinguish between good and bad directives A gives instead bears on how B can act while conforming to his reason to let A's directives robustly control his actions.

Finally, let me contrast the control-based account, not with Raz's specific discussion of authority, but with his more general discussion of normative powers. (Raz himself does not dwell on the relation between these two. But it is often assumed that authority is a normative power, and that an account of normative powers illuminates our thinking about authority.)

According to Raz, "a person's act is an exercise of a normative power if it brings about or prevents a normative change because it is, all things considered, desirable that that person should be able to bring the change about or prevent it by performing that act. Those who can exercise a normative power have a normative power to do so."³⁴ The account of justified authority I offered shares features with this characterization of (the exercise of) normative powers, in particular when it comes to the value (or, in Raz's terms, desirability) of A's having a certain power.

But they diverge centrally in the characterization of the power that is valuable or desirable. On the view I have offered, what is valuable is *not* (as for Raz) A's capacity to change B's normative situation. What is valuable is A's capacity to change – and, specifically, control – B's *actions*. It is because A's control over B's actions is valuable that A has control over what B ought to do, or the capacity to give binding directives to B.

Against proposals of the Razian sort, according to which what justifies the normative power is the value of someone's having the normative power, it is sometimes objected that they illicitly reason from 'it would be good if A's directives were authoritative for B' to 'A's directives are authoritative for B.' But no such worry arises with regard to the account I have offered, which moves from 'It would be good if A could control B's *actions* by mere say-so' via 'B has reason to contribute to A's having such control over his (B's) actions' to 'B has reason to let A's directives control what he (B) does' and (if this reason is undefeated, and perhaps other conditions met) 'A's directives settle what B ought to do.'

13.

Authority, it is often assumed, is distinctly difficult to justify. The answer I offered to the conceptual question helps shed light on why this is: A's directive is binding, or settles what B

³⁴ Joseph Raz, *The Roots of Normativity*, ed. Ulrike Heuer (Oxford: Oxford University Press, 2022), p.163.

ought to do. So an adequate justification of authority must explain the bindingness of authoritative directives.

This initial challenge is, however, effectively answered by the discussion of control and its value: because authoritative control requires that a directive robustly determines what B does, the value of control can explain why a directive has binding force for A. After all, only if a directive is binding does it give A the control that is so valuable. (This is compatible with the earlier observation that bindingness comes in degrees, and that, depending on the control that is worth realizing, directives may settle what is to be done more or less robustly and extensively.)

The discussion of control and its value identifies reasons for B to let A's decisions control his actions. Yet as the discussion of justificatory question shows, an account of *de jure* authority must also consider countervailing considerations. These include both considerations that defeat the reasons for letting A's decision control B's actions, and reasons that, though they do not defeat B's reasons to let A's decision control his actions, are also not defeated by those reasons and so leave B with the option of not letting A's decisions control his actions.

Some of these countervailing considerations are not special to authority. That A would direct B to do x, and doing x is incompatible with doing y, may be a reason against letting A's decisions control B's actions, if doing y is sufficiently important to B. But this is no different from cases *not* involving authority: in general, whether B should do x depends on whether doing x is incompatible with another action, doing y, and thus what the opportunity cost of doing x is.

Other considerations against A's controlling B's actions are, however, special to authority – not in the sense that they are limited to authority, but in the sense that the distinctive features associated with authority have a bearing on them.

A familiar worry triggered by authority is that it is incompatible with personal autonomy, or being the author of one's own life. If being the author of one's own life is a matter of giving shape to it through successive choices, then personal autonomy requires a range of diverse options from which the agent can choose. The bindingness of authoritative directives, which settle what the subject is to do, may be incompatible with the free choice among diverse options. So even if doing x would have been attractive by itself, being bound to do x may be problematic because it deprives the choice of doing x of its autonomous character. Such a cost in B's autonomy is a reason against B's letting A's decisions control his actions. But notice that authority's effect on B's autonomy depends on context. If what A directs B to do leaves plenty of meaningful choice for B (perhaps because it only forecloses a particular narrow set of options that would otherwise have been open), the effect on B's autonomy may be minimal or nil. If A's controlling B's actions is part of a broader practice that in turn helps make available new options for B, then A's authority over B, though it is pro tanto detrimental to B's autonomy, is all things considered beneficial for B's autonomy. And if being bound by A's directive is itself part of how B gives shape to his own life (because it is something to which he has consented, or because it is how he successfully pursues his personal projects and goals, as may often be true in instances of optional authority), then the benefits in self-authorship may once again outweigh any costs. So authority need not come at a cost in B's autonomy. But if it does come at such a cost, this is a reason against B's letting A's decisions control his actions that bears on whether A has *de jure* authority over B.

The observations about autonomy in the previous paragraph focused on B's being bound, but were independent of the fact that B's being bound was subject to A's control. (In other words: the same problems for B's self-authorship would arise if B had faced moral duties, or rationally non-optional choices, that were not created or imposed by someone else. B's autonomy is, for instance, limited by a natural disaster that leaves B with no rational choice about how to act to survive, or with no moral choice about how to act to protect others from harm.)

Other misgivings about authority have to do, however, with the interpersonal control that it involves. (The same is true with regard to coercion, compulsion, and manipulation.) A's controlling B's actions may be incompatible with a valuable relationship in which they could otherwise stand. (Friendship among equals, one may think, is incompatible with asymmetrical control by one over the other.) If so, then this too would be a reason against B's letting A's decisions control his actions. And in certain circumstances, A's giving binding directives to B may be insulting to B, because it shows a failure of respect for B's interests, in particular perhaps his interests in autonomy. This too would be a reason for B not to let A's decisions control his actions.

The discussion in the last few paragraphs has identified certain costs associated with authority relations. Yet interpersonal morality is not exhausted by an account of the interests that persons have. It is also constituted by deontic norms that govern how we may interact with one another, and how these interests count (or do not count) in justifying other-affecting actions. We have, for instance, personal rights that bar others from touching or using our body or property without our consent. These rights, to be clear, are not absolute. (If someone can save himself from drowning only by holding on to me, and his holding on to me only does minimal harm to me, then he may permissibly hold on to me even if I don't consent, and indeed object.) But they nonetheless serve to make certain resources – my body or property – asymmetrically available to me, for the pursuit of my personal projects and goals (or interests, for short), and bar others from using them unless sufficiently exceptional circumstances obtain.

These deontic norms play an important role in our thinking about freedom, understood not now as an achievement (leading an autonomous life) but as a *status* (being a free person, someone with certain entitlements vis-à-vis others). They capture the idea that we each have a life of our own to lead, not just in the sense that the choices we make and actions we undertake are descriptively ours (we are the ones making them, we are the ones who benefit when things go well and suffer when they go badly), but in a further normative sense: my choices and actions are asymmetrically available to me for the pursuit of my own personal projects and goals, just as your choices and actions are asymmetrically available to you, for the pursuit of your personal projects and goals. (Correlatively, an obvious way in which slavery is inherently wrong is that it treats the enslaved person's agency and body as asymmetrically available, not for the enslaved person's own personal projects and goals, but for pursuing the personal projects and goals of the person treated as the enslaved person's 'owner.') The asymmetric availability of my actions and choices limits what moral duties I have. (That you would benefit from my doing x more than I would benefit from not doing x is often not sufficient to put me under a duty to do x, if the choice whether to do x is one that is asymmetrically available for the pursuit of my personal projects and goals.) Similarly, it limits what others may do to me. (That you would benefit from my doing x more than I would benefit from not doing it is not sufficient to justify your compelling or coercing me to do x, if the choice whether to do x is asymmetrically available for the pursuit of my interests.)

And it explains why I have a complaint against certain kinds of treatments even if the actual burden to me is minimal or nil. (If an incoming student has a choice between a hundred majors, but then one, in which the student is anyway not very interested, is removed because the last faculty member retires, the effect on the student's interests may be minimal or nil. Yet if the major remains open, but the student is prevented from enrolling in it by another's threat that, if he so enrolled, his parents would be killed, the student has a very weighty moral complaint. Why? Because the coercer fails to respect that the choice whether to enroll has been fairly allocated to the student, to make in light of his own personal projects and goals. That entitlement reflects, not the student's weighty interest in enrolling or having the option of enrolling, but rather the importance of fairly distributing the resources – including opportunities for choice – useful for pursuing personal projects and goals.)

One might think that such deontic norms, and the corresponding status of B as a free person, bars A from having *de jure* authority over B in all but exceptional circumstances – effectively those circumstances in which B has either consented to such authority or another's interests are sufficiently urgent that B is under a moral duty to advance or protect them despite the asymmetric availability of B's action for the pursuit of B's interests. After all, authority gives A control over B's actions, yet B's actions are by assumption asymmetrically available to B, for the pursuit of B's personal projects and goals. (This is another way of explaining the temptation to insist the close connection between authority and moral duties.)

But this rests on a mistake. It is true that B's status as a free person limits when another's interests can impose moral duties on him, and thus when B is *morally required* to let A's decisions control his actions. But, as we observed earlier, *de jure* authority does not presuppose that B is subject to this moral requirement. It is enough that B has a unique undefeated reason to let A's decisions control his actions (i.e., is rationally but not morally required to do so), or that B has an undefeated reason to do so and in fact lets A's decisions control his actions. And, importantly, B's deontic status as a free person affects, in the first instance, what reasons *others* have, but not what reasons B has. If some value is best realized by letting A's decisions control my actions, and I have good, and uniquely undefeated, reason to do my part in realizing this value, then I have a uniquely undefeated reason to let A's decision control what I do, and A's decision controls what I ought to do. My status as a free person has no further bearing on the matter. Yet whether A has *de jure* authority over B is a matter of what reasons *B has*, not what reasons *others* have when it comes to dealing with B.

This is not to say, however, that B's deontic status as a free person is irrelevant for our thinking about authority. For an overall account of authority includes not just a theory of *de jure* authority (or when A's decisions are in fact binding for B in the requisite way), but also an account of what I earlier (in Section 1) called *legitimate* authority: an authority relation against which B has no moral complaint.³⁵ And that authority is justified is not sufficient to show that it is legitimate in

³⁵ Others too may have a moral complaint. A may have a complaint, for instance, if her ability to pursue her own personal projects and goals may be limited by the urgent and morally unavoidable task of making decisions for B. And a third party, C, may have a complaint if A makes decisions for B that fail to be suitably responsive to what B owes to C, yet nonetheless bind B. I flag these possibilities here without taking a stance on how they bear on the specific question of authority's legitimacy.

this sense, because (il)legitimacy is usually a matter of how *others* act in relation to B – and their actions vis-à-vis B are evidently governed by the deontic norms that constitute his free status.

First, it is possible that the justification of the authority relation rests on certain factual preconditions that wrong B. Consider the example (due to David Estlund) of the Dictator's Child.³⁶ In brief: A dictator has a child, The child sometimes gives directives to the dictator's ministers. When the ministers don't do what the child directs, the dictator gets angry and kills ten innocent people. The ministers have presumably a moral duty to prevent the death of ten innocent people (at least if they can do so at little cost to themselves), and so to do as the child directs them to do.

The child has *de jure* authority, in the sense that its directives control what the ministers ought to do. (Once the situation is such that the minister must choose between ignoring the child's directive and bringing about the death of ten innocents, the urgency of preventing the latter limits B's choices despite his status as a free person.) But the authority relation is morally illegitimate, because it rests on the wrongful disposition of the dictator, which creates an objectionable choice situation for the ministers that limits their freedom to respond to the child's directive in other ways. (That the dictator brings about this restriction on the minister's choice is objectionable given the minister's status as a free person.)

Second, the background conditions that justify A's authority over B may be unproblematic, yet A *abuses* the authority she has over B. Because B's actions are asymmetrically available for the pursuit of B's personal projects and goals, A, when giving directives to B, must be careful not to exercise her authority over B in the pursuit of ends to which B could object, because they are neither among B's own personal projects and goals nor among the moral duties that B owes to others despite his status as a free person. Imagine, for instance, that A's authority over B is justified by the moral importance of coordinating the activities of many citizens in the pursuit of the common good. If so, then A's directives must be based on considerations of the common good, as well as the interests of the subjects and how much they can be burdened for the sake of advancing the common good. If, by contrast, A's choice of directives is informed by the interests of others (her own, or those of third parties) which B is not morally required to advance or protect, and the advancement or protection of which is not among B's own personal projects and goals, then B has a moral complaint: A fails to treat his (B's) actions as properly his, in the normatively laden sense that these actions are asymmetrically available for the pursuit of his personal projects and goals rather than hers.

This does not mean that A's authority over B lacks justification. (Though, in extreme cases, it could: we have reason to escape abusive relationships, and so the fact that A abuses her authority over B is a reason for B not to let A's decisions control his actions. But that reason can be defeated by sufficiently urgent countervailing considerations.) But it does mean that A's authority over B is morally objectionable and thus illegitimate.

The distinction between *de jure* and legitimate authority, which falls out of the distinction between B's reasons and the reasons (and moral duties) of others in relation to B, helps address a complaint often raised against Raz's account of authority. The dictator's child example seems to

³⁶ David Estlund, *Democratic Authority: A Philosophical Framework* (Princeton: Princeton University Press, 2008), pp.xx.

meet the conditions of Raz's normal justification thesis, which he takes to establish 'legitimate authority.' Yet, the objection goes, that the dictator's child has legitimate authority is absurd.

What is absurd is that the child's authority is morally unproblematic. What is not absurd, I think, is that the child has *de jure* authority: the capacity to settle, by mere say-so, what the minister ought to do. Once we have at our disposal this distinction, the objection against Raz can be set aside.

14.

Because authority requires bindingness that can come apart from categoricity, authoritative directives need not be categorical reasons or moral duties. Still, one may wonder whether certain kinds of authority – perhaps not that of the coach or even the employer, but at least that of the state or the law, which is of special concern to political philosophers – requires categoricity, because only categoricity provides the kind of robust control over the actions of the subjects to which political or legal institutions aspire.

Consider first the authority that political and legal institutions claim over ordinary citizens (and, more generally, those within their jurisdiction). The law purports to authoritatively regulate a very wide range of behavior; and it may purport to have authority over an even wider range, insofar as it could regulate additional behavior that is currently left unregulated. The law does not take its authority to be optional for those in its jurisdiction, in the way that the coach's authority over the player is optional.³⁷

To this formal observation, about what the law claims, we should add a more substantive observation, about the practical role that we would want the law to play in our community. If, for instance, the law regulates social conflicts and disagreements, and its capacity to regulate them would be threatened if parties could simply opt out of its authority, then the law would lack the robust control necessary to provide certain valuable social benefits, such as the provision of peace and the fair resolution of disputes.

This explains why legal and political authority cannot be optional, in the way that the coach's or employer's authority is. But does it also follow that legal and political authority must be categorical in the further sense, that it (unlike parental authority) cannot be grounded in the interests of those subject to it?

The answer to that depends in effect on how robust authoritative control over a subject can be when it is ultimately grounded in the subject's own interests. As the example of parental authority shows, such non-categorical authority can in principle be very robust: even though parental authority is grounded in the interests of the child, the child cannot escape the parents' authority by revising personal projects and goals and thus changing the relevant interests. To show that political authority is inherently different, one would have to argue that, if political authority were grounded in the interests of those subject to it, then those subject to it would have the capacity to

³⁷ REF Raz, Green, Gardner?

revise their personal projects and goals and change the relevant interests in such a way that non-categorical political authority would lack appropriate robustness. Though I think adults do have more leeway to shape their own interests, I am quite doubtful that they have so much leeway that correlative non-categorical authority would lack the robustness required for it to fulfill important social goals. (To return to the earlier example: resolving conflict peacefully is sufficiently important for everyone that, simply by appeal to the subjects' own interests, the state can simply adopt a rule against opting out. It might, in other words, be in our interest that the control exercised over us be quite robust.)

Finally, let me flag that, though the authority that political and legal institutions claim over citizens is not optional, the authority it claims over its own officials very often is: public officials can, after all, withdraw from their job, and thus escape the authoritative demands that the institutions make on them *qua officials*. And in many contexts, authority over public officials matters as much for the justificatory role that it plays as for its restrictive role: that they have justifiably let someone else's decisions control their actions plays an important part in explaining how public officials can answer the complaints of private citizens who are burdened by their actions.³⁸ Given the centrality of public officials to political life, an adequate account of the role that authority plays in political and legal institutions must be attuned to this phenomenon.

³⁸ For insightful discussion of the room that officials have to engage in legitimate but unjust action pursuant to authoritative directives, see Jonathan Quong, *Political Philosophy: The Puzzle of Legitimate Injustice* (Princeton: Princeton University Press, forthcoming).

Bibliography:

- Adams, N.P. "In Defense of Exclusionary Reasons." *Philosophical Studies* 178 (2021): 235-53.
- Darwall, Stephen. "Authority and Reasons: Exclusionary and Second-Personal." *Ethics* 120 (2010): 257-78.
- Dyzenhaus, David. "How Hobbes Met the 'Hobbesian Challenge'." *The Modern Law Review* 72, no. 3 (2009): 488-506.
- Enoch, David. "Authority and Reason-Giving." *Philosophy and Phenomenological Research* 89, no. 2 (2014): 296-332.
- Estlund, David. *Democratic Authority: A Philosophical Framework*. Princeton: Princeton University Press, 2008.
- Frick, Johann, and Daniel Viehoff. "Authority without a Duty to Obey." *MIND* 132, no. 528 (2023): 942-51.
- Hobbes, Thomas. *Leviathan*. Edited by C. B. Macpherson. Harmondsworth, Middlesex: Penguin Books, 1986 [1651].
- Kelley, Mikayla. "A Control Theory of Action." *Australasian Journal of Philosophy* (forthcoming).
- Locke, John. "The Second Treatise." In *Two Treatises of Government and a Letter Concerning Toleration*, edited by Ian Shapiro. New Haven: Yale University Press, 2003 [1690].
- Parry, Jonathan, and Daniel Viehoff. "Instrumental Authority and Its Challenges: The Case of the Laws of War." *Ethics* 129 (2019): 548-75.
- Perry, Stephen. "Political Authority and Political Obligation." *Oxford Studies in Philosophy of Law* 2 (2013): 1-74.
- Quong, Jonathan. *Liberalism without Perfection*. Oxford: Oxford University Press, 2011.
- . *Political Philosophy: The Puzzle of Legitimate Injustice*. Princeton: Princeton University Press, forthcoming.
- Raz, Joseph. *The Morality of Freedom*. Oxford: Clarendon Press, 1986.
- . "The Problem of Authority: Revisiting the Service Conception." *Minnesota Law Review* 90, no. 4 (2006): 1003-44.
- . *The Roots of Normativity*. Edited by Ulrike Heuer. Oxford: Oxford University Press, 2022.
- Shepherd, Joshua. "The Contours of Control." *Philosophical Studies* 170 (2014): 395-411.
- Tadros, Victor. *To Do, to Die, to Reason Why: Individual Ethics in War*. Oxford: Oxford University Press, 2020.
- Viehoff, Daniel. "Procedure and Outcome in the Justification of Authority." *The Journal of Political Philosophy* 19, no. 2 (2011): 248-59.
- . "Service and the Moral Problem(S) of Authority." In *Engaging Raz*, edited by Andrei Marmor, Kimberley Brownlee and David Enoch, 487-509. Oxford: Oxford University Press, 2025.