Soft Law, Authoritative Advice, and Nonbinding Agreements

Andrei Marmor

One of the most fascinating developments in the domain of international law in the last few decades is the astonishing proliferation of nonbinding legal instruments, generally referred to as *soft law*. The forms and origins of various nonbinding legal instruments are quite varied, but what they have in common is the endorsement of norms, commitments, or directives, explicitly avoiding the imposition of legal obligations on the relevant parties. Needless to say, actual compliance with soft law varies considerably, but it is not negligible and sometimes surprisingly effective. International lawyers debate the legal status of soft law: Some doubt that it is genuinely legal material, others do not, and some suggest a middle ground between these two views. My concern in this essay, however, is quite different. I will not discuss the legal status of soft law. Instead, I want to focus on the rationale of nonbinding legal instruments from the vantage point of practical reason.

In spite of the multifarious occurrences of soft law in the international domain, we can draw a distinction between two main types: vertical and horizontal. Vertically, soft law comes in the form of nonbinding resolutions, directives, or quasi-judicial decisions, of various international authorities. Horizontally, soft law comes in the form of nonbinding international treaties (or parts of treaties that are nonbinding) and similar quasi-contractual arrangements between states.¹ My main purpose in this essay is to establish the rationale of *authoritative advice*, suggesting this rationale as a core model for thinking about vertical soft law. I will also try to show how we can expand a main

¹ Perhaps there are three kinds of soft law, actually. There are a number of global initiatives, some entirely set up and maintained by NGO's, like *Social Accountability International*, others, like *Global Impact*, under UN auspices, that strive to set international norms and monitoring standards for socially responsible global corporate conduct. Many international corporations sign up for these initiatives, though invariably on a voluntary and self-regulating basis. Critics of these initiatives worry that this mechanism, gaining some momentum in recent years, serves the interest of global corporations by vindicating their claim to self-regulation, thus helping them avoid legal accountability. Others think that given the unfeasibility of more robust legal regulation, voluntary self-regulating mechanisms may achieve some good. Either way, these international standards are, at best, a borderline case of soft law. I do not purport my analysis here to apply to this third category, however.

aspect of this model to account for nonbinding agreements, thus explaining the rationale of horizontal soft law.

The idea of an authoritative advice might strike us as oxymoronic: One would think that it is an essential aspect of practical authorities (perhaps in contrast to epistemic authorities) that they purport to impose obligations. If a directive is authoritative, it purports to be binding or obligatory. And if a directive is explicitly not meant to be binding or obligatory, it is not really authoritative. But this is precisely the dogma that I aim to challenge here. I want to show that an authoritative advice is not an oxymoron. And then I will try to show that a similar rationale, with an additional element of mutual accountability, applies to the idea of nonbinding agreements. Taken together, I hope that this will give us some tools to account for the practical reasoning underlying the idea of soft law.

One more clarification before we delve into the issues. My purpose in this essay is purely philosophical, that is, to give an account of the practical reasons that are instantiated in the type of cases I discuss. I do not claim that such an account can provide answers to the empirical questions of why states and other international actors actually endorse nonbinding instruments. There is a growing literature on these issues, purporting to explain why it is in the interest of various international agents to devise nonbinding instruments, what are the conditions that account for different levels of compliance with soft law, and similar issues.² I do not purport to add anything to this legal, economic, and political science literature. Nor do I intend my arguments to depend on any particular behavioral explanation of this phenomenon. My aim is only to suggest one, rather abstract, way of thinking about soft law from the perspective of the practical reasons in play.

A. Authoritative Advice & Vertical Soft Law

1.

Philosophical theories about the nature of practical authorities and the conditions of their moral legitimacy vary considerably. I have attempted my own stab at this issue

² See, for example, D. Shelton, (ed.), *Commitment and Compliance: the role of nonbinding norms in the international legal system* (Oxford, 2000, 2007).

elsewhere.³ But I think that we have all taken it for granted that it is an essential feature of practical authorities that they purport to be binding or obligatory. The assumption is that if A is an authority *vis a vis* B in matters of type C, then A's authoritative directive to B to ϕ in C necessarily purports to impose an obligation on B to ϕ in C. And if the authority of A over B in matters C is morally or otherwise legitimate, then the assumption is that B has an obligation to ϕ in C, and at least partly because ϕ -ing has been directed by A. The obligation may be *pro tanto*, or it may take some other form, but an obligation it is. Perhaps the obligation is not necessarily a moral one, it may be institutional or legal in nature, but again, some obligation is necessarily in play. Now of course this is a very simplified way of putting things. There are many kinds of directives authorities can issue besides directly imposing an obligation; they can grant or withhold various rights or permissions, grant or withhold normative powers, etc.,. But all these normative changes can be reduced to obligations imposed on someone. When an authority grants X the right to ϕ , for example, *ipso facto* it imposes an obligation on some other agent(s), Y, to enable X to ϕ in some relevant sense. When the authority grants X the power to ϕ *vis a vis* Y, *ipso facto* it imposes an obligation on Y to comply with X's exercise of the power. And so on and so forth. Normative changes introduced by authoritative directives are reducible to obligations.⁴ This is what is generally meant by the idea that the directives of a practical authority purport to be binding. And then we assume that whether particular authoritative directives are actually binding or not would depend on the authority's legitimacy *vis a vis* the putative subjects in the relevant context.⁵ I do not aim to suggest that this picture is generally wrong or misguided. I think that, by and large, it is correct. But I would like to show that there is, nevertheless, a clear sense in which a suggestion or an advice, not meant to be obligatory, can be meaningfully authoritative.

³ "An Institutional Conception of Authority", 39 *Philosophy & Public Affairs* (2011), 238

⁴ See, for example, Raz "The Problem of Authority: Revisiting the Service Conception" in his *Between Authority and Interpretation*, (Oxford 2009), ch 5.

⁵ We probably owe this general framework of how to think about practical authorities to J. Raz's influential work in *The Morality of Freedom*, (Oxford, 1986) ch 2-4. Though many different accounts of practical authorities have come up in the literature since then (including mine, if I may say so), some of them critical of key components in Raz's theory, Raz's basic ideas remained very influential. One notable exception is Stephen Darwall's recent work on the second personal standpoint, claiming that it gives us a very different conception of practical authority (see, for example, S. Darwall "Authority and Reasons: Exclusionary and Second Personal", 10 *Ethics*, (2010), 257.) . It is very doubtful, however, that Darwall's rather peculiar understanding of what counts as an authority can capture our everyday notion of practical authority. See my "The Dilemma of Authority", 2 *Jurisprudence* (2011), 121, at 134-7.

Let me start with a few mundane examples. Consider J, a young assistant professor who happens to be a tad sluggish with his research, spending a lot of his time and energy on all sorts of other things instead. Now suppose that J gets the exact same advice, let us assume, expressed identically, from two different sources: His best friend, and his department chair. Both advise him (separately and unrelatedly) to pay much more attention to his research, expressing concern that he might fail to get tenure otherwise. One would be a friendly advice, the other, I suggest, might well be an authoritative advice. When the department chair advises her young colleague to pay more attention to research, her expression is not meant to impose an obligation on J. And yet, it is meant to make a practical difference, different from the advice of a friend or a stranger. The department chair speaks here, I will assume, not as a friend (which she may or may not be as well) but as the department chair. And it is meant to be significant that her advice is issued in that official capacity. At least we can easily postulate such a context.

A word of caution is in place here. We need to acknowledge that it is sometimes very difficult to distinguish between a genuine advice or suggestion, and some veiled threat or an implicit obligatory directive. Generally, the content actually conveyed by such expressions is a matter of pragmatics; the same sentence uttered by the department chair may intend to convey different contents, possibly implicating more than advice. But it does not have to, and for simplicity's sake, I will henceforth assume a condition of full sincerity. That is, I will assume that in all the cases I discuss, when something is expressed as an advice, the content actually intended by the speaker and conveyed by her is, indeed, nothing more than an advice (and nothing less). Furthermore, I am assuming here that the differences between advice, friendly suggestion, encouragement to do something the other has reason to do, and things like that, are often rather subtle. One should therefore take the notion of advice, as I use it here, expansively, encompassing similar concepts in the vicinity.

Let me give two more examples of what I take to be authoritative advice. Continuing our story of J, the assistant professor, suppose that he is now concerned about his tenure prospects. So he looks up the university guidelines for tenure in his field, and what he finds, among other directives, is a guideline saying that “a published book manuscript written during the period is normally required.” Slightly puzzled by this

formulation, J reassures himself, however, that a book publication is not mandatory, only advised. Thus he concludes that he is not under an obligation to publish a book in order to qualify for tenure. Well, he is right about that, but the university guideline about the book requirement is not meant to be pointless or vacuous either. It is, again, what I take to be an authoritative advice.

You might worry that the examples about J pertain to reasons for action that are agent-centered. After all, J does not owe it to anyone to get tenure, it something that is entirely up to him and his own preferences. I am not sure why would it matter for our purposes here⁶, but we can give other examples. The employment context provides many cases. Thus suppose that a big law firm issues a series of guidelines to its employees. Many of them would be mandatory requirements, but some might take the form of authoritative advice. The firm may advise, but explicitly not require, its legal associates to wear “business attire” at work; it may advise its employees to be particularly mindful of showing respect to women and racial minority clients; it may advise but not require certain methods of communication with business partners; and things like that. Or think about a large employer mindful of increasing the diversity of its workforce, adopting a moderate policy of affirmative action by explicitly advising (but, crucially, not requiring) its mid-level managers to consider a preference for minority job applicants if all other things are equal. The managers would rightly consider this policy as advisory, not mandatory. The firm does not obligate them to tilt the scale in favor of minority candidates whenever possible, it only advises them to do so.

In order to show that these are examples of authoritative advice, we need to show that the subjects’ reasons for action depend on the *say so* of another. We cannot begin to understand the nature of practical authorities without bearing in mind that it is one of those cases in which it matters that a particular expression has been actually expressed by someone, and it matters who the speaker is. Spectators at a soccer game can shout at one of the players “out” as long as they want, it would not give the player any authoritative or

⁶ It might matter in the following sense: There might be a way to construe the example in terms of epistemic authority, not practical authority. One could claim that the department chair’s official advice to J, and the university guidelines about tenure requirements, are just a kind of expert advice, carrying some epistemic authority, but not practical authority. But this seems to me like a stretch. It is, of course, possible to stipulate that the department chair’s advice is given to J by way of making an expert recommendation, just informing J of what is required of him to get tenure. But this changes my example.

institutional reason to leave the field. But if the umpire says “out”, the player is, we would assume, under an obligation to leave. Reasons for complying with an authoritative directive crucially depend on the say so of the authority, and thus it matters who says so. Perhaps the advice to J to focus on his research is sound on its merits, he really needs to do that. But it matters, or so I shall argue, that in one case, the advice is given by a friend, and in another, by the department chair in her official capacity as the department chair. There is no authoritative directive without a speech act, and it matters whose speech act it is and in what capacity it is expressed. Elsewhere I have called these kinds of reasons for action *identity related*.⁷

Pointing out, however, that reasons for action are identity related is not enough to establish that the speech act in question is authoritative. We find identity related reasons for action in other contexts as well, such as requests of a friend or expert advice. If we want to explain what makes authoritative advice authoritative, as it were, we need more; we need to explain how a piece of advice expressed by someone in a position of authority makes a difference, and what kind of difference, to the reasons for action of their putative subject. If it is not meant to impose an obligation, what else can it do?

The difficulty here is compounded by the nature of advice. It is, quite generally, the nature of advice that it is not meant to create new reasons for action, so to speak, but only to remind the addressee, or point out to her, reasons for action she has anyway. When the friend tells J that he should focus on his research, for example, the friend does not intent to give J reasons for action he had not had prior to the advice. On the contrary, the whole point of the advice is to bring to J’s attention reasons for action that apply to him anyway, regardless of the speech act of advising. And this is quite generally how we would think about the nature of advice (or well intentioned suggestions and things like that); they aim to point out reasons for action that apply to the addressee anyway, regardless of the advice. So the crucial question we must answer here is why would it matter to the nature of an advice that it has been issued by someone in a position of authority? How can it make a difference to the kind of reasons for action in play?

⁷ See my *Philosophy of Law*, (Princeton, 2011), at 62-3.

2.

I think that the reasons for action generated by an authoritative advice resemble the kind of reasons for action we find in rebuttable presumptions in the law. Presumptions in law, at least in one familiar type of cases, operate as devices for shifting the burden on the subject to come up with defeating reasons that are clear and convincing enough to defeat the first order reasons the presumption stipulates.⁸ Suppose, for example, that there is a legal presumption that if condition C applies to subject A, the law assumes that A is X. If the presumption is of the rebuttable kind (as most are), the practical effect of the presumption is to shift the burden on to A to show that despite the fact that condition C applies, A is not X. Presumptions, we can say, involve a first order reason for assuming, legally speaking, that some state of affairs applies to A, and a disjunctive reason shifting the burden on to A to disprove the assumption if it is not correct.⁹ I submit that a very similar rationale applies to the kind of reasons for action authoritative advice creates.

For lack of a better idea, I will call these kind of reasons presumptive reasons for action. R is a *presumptive reason* for action for A to φ in C, *iff* --

R is a defeasible, first order reason for A to φ in C, *and*
if R is allegedly defeated by other reasons that apply to A in C, there is an expectation *from A* to demonstrate that the defeating reason is present and compelling, as such, under the circumstances.

Simply put, the idea is that presumptive reasons point out to the subject that there are reasons for action that apply to her, acknowledging (explicitly or implicitly) that the reasons are potentially defeasible, but expecting the subject, if she relies on a defeating reason, to prove or demonstrate that the first order reasons are indeed defeated in her case. Notice that presumptive reasons are prevalent in many ordinary contexts. We often expect someone to do something or else explain why they failed to do it. If you are

⁸ Rebuttable presumptions in law are usually not about reasons for action, at least not directly, they tend to stipulate a certain state of affairs conditional upon a given set of circumstances. Though of course they may affect reasons for action indirectly.

⁹ Some jurisdictions have elaborate rules about these things; sometimes the burden shifted to the subject is only the burden to adduce the relevant evidence, sometimes it is the heavier burden of proof. These technicalities do not matter for our purposes, the legal example is only illustrative.

invited to a meeting or a dinner party or such, you are expected to be there on time. And if you are late, you are expected to explain. Perhaps you had a good reason to be late, but you may need to tell your colleagues or hosts what the reason is.

I am using the word “expectation” advisedly here, and purposefully hedging a bit. The expectation to demonstrate that the first order reasons are defeated is not necessarily or even typically an obligation. Rather, it might just be something that would be deemed appropriate, and perhaps reasonably asked of the subject, or even demanded of her, even if it does not have a binding or obligatory character. You are not obligated, for example, to explain to your colleagues why you were late for the meeting; but it is something that they can rightfully expect you to do. As I hope it will become clear as we go along, different contexts might justify different kinds of expectations in the relevant sense.

Authoritative advice gives the subject presumptive reasons for action. When advised by an authority to φ , the subject would be right to treat the reason to φ as defeasible, but it is up to the subject to demonstrate to the authority that the relevant reasons are, actually, defeated in her case, if they are.¹⁰ Think about the example of the university tenure guidelines strongly recommending, but not obligating, candidates to have published a book. What is the point of it? I suggest that it is precisely an example of presumptive reasons; the university would assume that you need to have published a book to get tenure, and yet it acknowledges that the book requirement is defeasible. Perhaps if you have a sufficient number of high quality articles, the lack of a published book would not count against you, at least not decisively. But you will need to convince the authorities that that is the case; you will need to make a point of overcoming your non-compliance with strong and clear defeating considerations. And I think that the same kind of reasoning applies to the other examples. When the employer advises, but does not obligate, the employees to do this or refrain from doing that, it aims to point out to the employees that they have first order reasons for action, acknowledging that the reasons

¹⁰ I am well aware of the fact that mandatory directives are defeasible as well. The subject’s obligation to φ can be defeated under various circumstances; but then we would normally assume that the defeater has to be either a conflicting obligation or else the kind of reason that is not excluded as potential defeater by the obligation in question. In other words, without going into any details about the nature of obligations, we can acknowledge that obligations are defeasible, but there are various constraints on what can count as a justified defeat. I am assuming here that there are no such constraints on justified defeat of ordinary first order reasons for action involved in presumptive reasons for action. The only additional constraint is to explain the defeat.

can be defeated. But then it is up to the employee to demonstrate that this is the case. And similarly, when the department chair advises J to focus on his research, a presumptive reason is in play. Perhaps J can pull it off without putting in the extra work, perhaps he knows something that the department chair is not aware of. But the point of the authoritative advice is to put J on notice that it will be up to him to show that the relevant reasons, in his case, are actually defeated, if they are. This, quasi-procedural aspect is precisely what distinguishes the official, authoritative, advice of the department chair from the identical advice of J's friend.

If you are familiar with Joseph Raz's theory of practical authority, you might wonder what is the difference between presumptive reasons, and Raz's account of the *preemption thesis*.¹¹ So let me try to explain briefly. The preemption thesis consists of two elements: Raz argued that a legitimate authority's directives create both *preemptive reasons* for action and *protective reasons*. The *preemption* consists in the idea that when an authority requires the performance of an action, it gives a reason for its performance which is not to be added to the subject's balance of reasons when assessing what to do, but should *replace* some of them. If my account above is correct, it follows that authoritative advice is not preemptive in nature. Pointing out the first order reasons for action that apply to the subjects is not meant to replace any of the subjects' assessment of the first order reasons for action; but it is not meant to be added to them either. Instead, the authoritative advice is there to bring certain reasons that apply to the subject's attention, as it were, making them salient.

According to Raz, however, authoritative directives also constitute *protected reasons*: A protected reason to ϕ is a "first order reason to ϕ and an exclusionary reason not to fail to ϕ for a certain range of excluded reasons".¹² So here is the main difference between protected reasons and presumptive reasons: Protected reasons are *exclusionary*, they purport to exclude from practical consideration certain types of reasons for counting

¹¹ Raz's formulations of the preemption thesis, and the terminology he used, have not been entirely consistent over the years. The main articulation of his theory of practical authority is in his *Morality of Freedom*, ch 2-4. Further refinements and slight modifications of the thesis are found in his *Ethics in the Public Domain*, ch 9, *Between Authority and Interpretation*, ch 5, and most recently in "On Respect, Authority and Neutrality: A Response", *120 Ethics* (2010), 279.

¹² Raz, *Ethics in the Public Domain*, (Oxford, 1994) at 198. I tried to explain Raz's arguments and formulations in greater detail in my "The Dilemma of Authority", 124-9.

as potentially defeating the first order reasons for action. For example, when mother directs teenager son to be home by 7:00 for dinner, she purports to give son a first order reason to be home for dinner, and second order reasons to exclude from potentially defeating reasons a certain range of possible defeaters, such as an opportunity to stay longer in town to hang around with friends, and things like that. Not all potential defeaters are excluded, of course. If an accident happens and son is delayed by the need to help a friend, that would OK, such defeaters are not excluded.¹³

The disjunctive reasons constituted by presumptive reasons, on the other hand, are not of the exclusionary type. Their point is not to exclude from practical consideration certain types of reasons from counting as potentially defeating reasons. On the contrary, authoritative advice typically involves an implicit recognition that nothing is excluded by way of potentially defeating reasons. It only shifts the burden to the subject of the advice of proving or demonstrating that a sound defeating reason applies under the circumstances, if it does.

3.

We should still wonder, however, what goes wrong if the subject of an authoritative advice fails to treat it as a presumptive reason, and simply ignores the advice? After all, we often ignore advice that we get, and in most cases, there is nothing wrong with it. You may have advised me to learn how to ski, for example, and perhaps you were right, it would be great fun. But if I fail to heed your advice, there is really nothing that I need to do about it; certainly I don't have to explain myself to you or to justify my decision to ignore your advice.¹⁴ Why should it be different with authoritative advice?

Allow me two preliminary points before I try to answer this question. First, let me make an assumption that I have relied upon here more explicit. My assumption is that we can talk about an authoritative advice only on the condition that some kind of an

¹³ As I explained elsewhere in some detail, protective reasons are quite ubiquitous and do not necessarily involve obligations. (Nor has Raz ever claimed otherwise.) See my "The Dilemma of Authority" at 125-6.

¹⁴ In some contexts there might be a justified expectation to show the advisor that you have taken their advice seriously, that you thought about it, or something like that. But this is not something that generally follows from the nature of an advice; it has something to do with the relationship between the relevant parties (whether they are friends, colleagues, strangers etc.) and other contextual factors.

authoritative relation between the relevant parties is already in place. No advice or suggestion or anything similar could possibly constitute an authoritative context.¹⁵ The context has to be such that A is already in an authoritative position *vis a vis* B with respect to the relevant issues, and then the question is whether in this authoritative capacity, A's directive to B can take the form of an advice, as opposed to an obligatory directive. Admittedly, this presupposes that we can articulate the conditions for what it is for someone to have practical authority *vis a vis* another, regardless of the particular contents (and/or merits) of the putative authority's directives. In other words, we must be able to establish what it is to have practical authority regardless of what would make an authority legitimate or otherwise warranted. This thesis is not uncontroversial, but I have argued for it at length elsewhere, and I will not repeat the argument here.¹⁶ I just want to make it explicit that I rely on this assumption.

As another preliminary, let us return to the advice of a friend. As I mentioned earlier, advice, by its very nature, purports to point out reasons for action that apply to the addressee anyway, regardless of the advice. But then a question arises: What could be the difference between advice you get from a friend, and the same advice that you might get from a complete stranger? Is there a general principle that friendship binds you to some *pro tanto* expectation of compliance? Can friends expect that their advice be heeded partly because they are your friend? That I seriously doubt. It seems to me that an expectation of heeding the advice, *pro tanto* and defeasible as it may be, is just the wrong kind of expectation in the context of friendship. One who has such expectations would seem to misconstrue the values of friendship.¹⁷ Friendship, we should think, does not give you authoritative powers over your friend. You don't get to decide things for your friend based on your friendship, only to care for them, to offer your support, your help, your wisdom, or what have you. The difference, if there is one, between advice offered by a friend and one offered by a stranger pertains to presumed motivation. When you get an

¹⁵ This may be different with epistemic authorities. My discussion here is confined to practical authorities.

¹⁶ "An Institutional Conception of Authority". Raz (in *The Morality of Freedom*, ch 2), for example, expressed doubts about this direction of the argument, suggesting that we cannot have an account of what it is to have practical authority without first understating the conditions of an authority's legitimacy. I am not sure whether he still holds this view.

¹⁷ I am not denying, of course, that there are various contexts in which friendship generates reasons for action; a friends' request would normally give you a reason to comply, sometimes perhaps even an obligation to do so. And this may well derive from the relationship of friendship.

advice from your friend, you would normally assume that the friend's motives are driven by concern with our own well being, something that you may not be so sure about when the advice comes from a stranger. And there may be other, similar factors in play. But the one thing friends, I submit, cannot legitimately expect is that their friendship endows the advice they give with a reason, that derives from the friendship, to heed the advice.

It is precisely in this respect that authoritative advice differs from advice of friends or strangers. In the relations between legitimate authorities and their subjects there is a warranted expectation of compliance. And this expectation is closely entangled with a notion of trust. In order to function properly as an authority in a certain domain, the authority needs to have some level of confidence that their subjects would normally comply with their directives, and the subjects would need to have some level of confidence that other subjects comply. In other words, authorities (assuming that they are by and large legitimate) and subjects need to trust each other. Of course a practical authority can gain confidence in compliance without trust; the authority can intimidate or terrorize subjects into to compliance, which sadly, often happens. But I am assuming that this is not the kind of authority that we would deem legitimate and worthy of the putative subjects' respect. The proper functioning of a legitimate authority requires mutual trust. Without a decent level of trust in general compliance with an authority's directives, constant and relentless monitoring of compliance, and punitive reaction to non-compliance, would have to become the norm, as it were. But incessant monitoring of compliance is very unhealthy; authorities would find it very difficult to fulfill their functions if they need to monitor compliance on every little step, and subjects would find it intrusive, perhaps even despotic, if the relevant authority exhibits no general trust in their compliance and scrutinizes it on every turn.

Now I think that we have the answer to the question I posed above: what goes wrong with flouting an authoritative advice (assuming, of course, that the authority is legitimate and its advice is well within the scope of its warranted authority)? The answer is that it betrays the kind of trust that should obtain in the relations between authorities and their subjects. Unlike friends, legitimate authorities should be able to expect that subjects comply with their directives, and that, I think, includes directives that aim to generate presumptive reasons for action. There is no reason to think that the level of trust

that authorities need to have in their subjects' willingness and tendency to comply is affected by the kind of reasons for action they direct. Whether the reasons they purport to give by their authoritative directive is obligatory, or only presumptive, does not bear on the desirable level of confidence, and mutual trust, that the authority and the other subjects can justifiably expect about compliance.

Before we move on, let me answer a possible objection or, rather, dispute an alternative explanation. Several philosophers have argued over the years that it is an essential aspect of authority–subject relations that the subject is *accountable to the authority* for their compliance with the authoritative directive. Thus, the idea is that an obligation to obey an authority's directive is directional, one that is *owed by the subject to the authority* in question.¹⁸ Why is this relevant here? If the accountability view of authority-subjects relations is correct then we may not need the idea of trust to explain what goes wrong when a subject ignores a piece of authoritative advice. We could simply rely on the accountability principle instead.¹⁹

Tempting as this view is, I doubt that it works. Proponents of the accountability view allege that whenever a subject fails to comply with a directive of a legitimate practical authority, the subject thereby *wrongs the authority*; you do or fail to do something that the authority can rightfully resent and hold you accountable for it. Elsewhere I expressed in some detail my doubts about this directionality of the obligation to comply with practical authorities.²⁰ On my view, obligations to comply with a legitimate authority's directives are typically obligations owed to the members of the practice or the institution on whose behalf the authority operates, not a directional obligation owed to the authority in question. I will not repeat those arguments here, only

¹⁸ See, for example, M. Gilbert, *A Theory of Political Obligation*, at 245-8, S. Darwall "Authority and Reasons: Exclusionary and Second Personal", *120 Ethics*, (2010), 257, S. Hershovitz, "The Authority of Law" (Routledge Companion).

¹⁹ With some modifications, of course, because advice does not impose an obligation.

²⁰ "The Dilemma of Authority", 135-7, and "An Institutional Conception of Authority", 255-50.

point out that the accountability view is even more problematic when applied to an explanation of what goes wrong with flouting an authoritative advice.²¹

The problem for the accountability account in cases of authoritative advice stems from the nature of advice, not (only) from the nature of practical authorities in general. One just does not wrong the person who gives you an advice by ignoring it. An exceptional case might prove the point here. Parents often feel resentment or personally disappointed when their child ignores an important advice they had given the child. This is quite natural and perhaps even warranted in many cases. But notice the importance of the special parental relations that makes such feelings of resentment or disappointment potentially appropriate: it would be rather strange to assume that what makes them appropriate is the authoritative role that parents occupy *vis a vis* their children (which they normally also do, of course). The potential appropriateness of disappointment or even resentment in such cases stems from the special bond and very personal and affectionate relations that are normally present in parent-child relationships. The department chair, however, or the manager at your workplace, is not your parent, and it would be wrong for them to expect to be treated like one. In short, I do not think that recalcitrant subjects wrong the authority whose official advice they ignore. But authorities, and other subjects of the authority, do need to share a certain level of trust in each other. Authorities need to be able to assume that if they give you a presumptive reason for action, you will either comply or else demonstrate that you have a good reason not to comply. They need to assume this in order to fulfill their functions as practical authorities, not because you would wrong them by noncompliance. By noncompliance you wrong the proper functioning of the institution that needs the authoritative structure it has, and those who normally benefit from it.

²¹ I should note that the standard accountability view would face some serious difficulties in the context of international law, quite generally. The subjects of authoritative resolutions in the international domain are states, and it is far from clear that states owe their duties of compliance exclusively or even mainly to the authorities in question. States, and the agents acting on their behalf, may owe their compliance obligations to their own citizens as well and perhaps, in some cases, to citizens of other states. The directionality view of duties to comply with authorities does not sit easily with the complexities of international law. Accountability, as will see in the next section, however, is of crucial importance to the rationale of nonbinding agreements.

4.

Let us move on to the next stage: assuming that there is such a thing as authoritative advice, would it help us to account for the rationale of vertical soft law, such as nonbinding resolutions of authoritative international institutions? Notice that our sole concern here is vertical soft law; I am not claiming that horizontal soft law is a form of authoritative advice. In order to explain the latter, we need more groundwork, that will be presented in the next section. At this stage, however, I want to suggest that the nonbinding directives of international authorities can be seen in light of the idea of authoritative advice. Remember that I have explicitly acknowledged earlier that the very idea of an authoritative advice only makes sense on the assumption that an authoritative relation between the relevant parties is already in place. Only someone who is already in an authoritative position *vis a vis* another can issue authoritative advice to the relevant subject. So this model does not directly apply to horizontal soft law constituted by international treaties. But it does apply to directives of authoritative institutions, including those previously established by treaties, as most international institutions are.

Consider, for example, the UN General Assembly resolutions adopted in the early 1990s (in several stages) concerning the eradication of driftnet fishing.²² The UNGA Resolutions were nonbinding, and yet, for various reasons that we need not elaborate here, surprisingly effective. Compliance, however, is not our concern here. I am using this example to demonstrate the practical reasoning that is involved in the issuing of nonbinding resolutions by an international authoritative institution, suggesting it as a fairly standard example of authoritative advice. The point of authoritatively advising the cessation of driftnet fishing was to make the reasons against this harmful practice as salient as possible, and to shift the burden of demonstrating defeating reasons on states who wish to continue the practice despite its tremendous harmful effects. Now you might think, and rightly so, that the main purpose of the UNGA resolution was to bring about an actual change in the practice, not to advise. Of course that is the main point, it's the main point of any directive, whether obligatory or advisory. The question is not why regulate, but why regulate in the form of a nonbinding resolution. The answer, no doubt,

²² See D. R. Rothwell, "The General Assembly Ban on Driftnet Fishing", in *Commitment and Compliance*, 121-146.

is often the simple one, that the relevant authority cannot accomplish, legally or politically, a resolution that would be legally binding. Often, no doubt, soft law is a second best solution. Be this as it may, the question we aim to answer here is about the rationale of the directive when it is, for whatever reason, explicitly nonbinding. And I submit that in these kinds of cases, the idea of an authoritative advice provides a pretty good model, at least in terms of the kind of practical reasons involved.

By issuing the nonbinding resolutions to eradicate driftnet fishing, the UNGA aimed to accomplish two things: First, to make it public and salient that driftnet fishing is extremely harmful and better be discontinued. It is noteworthy that the resolution was backed up by extensive scientific research, demonstrating the harms of driftnet fishing and mandating continuous scientific monitoring.²³ It is also noteworthy that the original resolution called upon the Secretary-General to bring the resolution to the attention of the international community. Making the issue salient was explicitly part of the point of the resolution. Second, to generate an expectation or a warranted demand on noncompliant states to come up with very convincing reasons, if there are any, to continue with the practice in spite of the many reasons to the contrary. The original resolution made it explicit that the reasons against driftnet fishing in the high seas are defeasible, saying that “such measure will not be imposed in a region should effective conservation and management measures be taken....”, but also shifting the burden of proving the defeating reasons on noncomplying states, requiring defeaters to be based on “sound analysis to be jointly made by concerned parties of the international community....”.²⁴ In short, the structure of presumptive reasons here is expressed very clearly. And indeed, when noncomplying states failed to demonstrate sound defeating reasons, pressure mounted on them to invest in alternatives. Many factors contributed to the success of this project, not the least considerable US investment in it, both pressuring and aiding other states in finding viable alternatives.²⁵ Once again, however, the success story here only demonstrates that authoritative advice, even in the international domain, sometimes

²³ This is quite typical of similar resolutions, e.g. IAEA recommendations to protect nuclear material, or the recommended measures by the ATCM regarding the Antarctic.

²⁴ (44/225 1989)

²⁵ See Rothwell, “The General Assembly Ban on Driftnet Fishing”, in *Commitment and Compliance*, 121-146.

achieves its purpose. No doubt, it often fails. But then remember that binding resolutions in international law often fail as well.

B. Nonbinding Agreements & Horizontal Soft Law.

5.

Even if the idea of authoritative advice provides a model for thinking about vertical soft law, the practical reasoning behind nonbinding international treaties remains to be explained. International treaties are agreements between states, not authoritative directives. But the puzzle about nonbinding treaties is very similar to that of nonbinding authoritative directives, both have a similar air of paradox to them. Contractual agreements normally consist of mutually undertaken obligations by the parties to the agreement; so what would be the point of reaching an agreement that is explicitly nonbinding? The idea that I would like to articulate here is twofold: first, that nonbinding agreements also tend to establish presumptive reasons for action, very much like authoritative advice, in this respect. Additionally, however, agreements constitute accountability relations between the parties, rendering them answerable to each other. In other words, nonbinding agreements establish presumptive reasons for the parties, coupled with mutual accountability relations that enable those presumptive reasons to play their practical role.

Let us work with a simple model: Suppose that two states, A and B, sign a treaty that contains a commitment, applying to both parties, to φ in context C. However, let us stipulate that they make the agreement explicitly nonbinding; the undertaken commitment to φ in C is explicitly non-obligatory.

The main question that arises here is what is it, exactly, that the parties have actually committed themselves to? What kind of reasons for action such a nonbinding agreement can generate? A very tempting answer might be the following: a binding agreement is the undertaking of an obligation to φ . A nonbinding agreement is the undertaking of an obligation *to try* to φ . So we might think that the nonbinding agreement between the parties is essentially a commitment to try, not a commitment to do or to accomplish. It is as if each party is saying to the other: “I cannot commit myself to doing

φ , but I commit myself to trying, I will attempt to φ .” This seems neat, but wrongheaded. Generally speaking, a commitment or a promise to try to φ would be utterly insincere if the agent undertaking it does not actually have the intention to φ . An attempt to do something requires the intention to do it.²⁶ There are some exceptions, perhaps. For example, it seems possible that one would commit to trying without having the intention of succeeding when they know that they are bound to fail, and the attempt is carried out only to demonstrate that they will fail. But this is not quite relevant to nonbinding agreements. Surely parties to such an agreement do not make it only to show to the other party that they are bound to fail. That does not make much sense. Perhaps there are also some cases, rather exceptional however, in which it makes sense for an agent to try to φ , without intending to succeed, hoping that by trying to φ they can achieve something else in the vicinity that they intend to accomplish.²⁷ But again, if this is really a case in which one can try to φ without intending to φ , the case would be a very unusual one, and hardly relevant to the rationale of nonbinding agreements. In short, the construal of nonbinding agreements in terms of mutual undertaking of commitments *to try* without a commitment to do or to accomplish, does not make much sense. It inevitably renders such agreements profoundly insincere.

Thus the question remains, what kind of reasons for action might be generated by a nonbinding commitment of A and B to φ in C? If it is not a commitment to try, what is it? I would suggest that it is a combination of two elements: the commitment consists in recognizing and affirming a presumptive reason for φ -ing in C, accompanied by establishing an accountability relation with the other party, subjecting each other, as it were, to criticism, and perhaps even to some pressure, if the presumptive reasons are not complied with. In other words, it is a mutual commitment either to φ in C or to demonstrate to the other party that reasons for φ -ing in C are clearly defeated, if they are.

Before I try to explain this in greater detail, let me limit the scope of the argument here. International treaties, binding and nonbinding alike, are signed by states for a great variety of reasons, political, economic, military, scientific and whatnot. If the international treaty is of the binding type, the reasons for joining the agreement should

²⁶ See, for example, G. Yaffe, *Attempts*, ch 2; M. Bratman, *Intention, Plans, and Practical Reason* (1987)

²⁷ Bratman, (1987)...

explain why it makes sense for the signatory states to undertake the relevant obligation or, at least, to appear to be doing so. These types of cases are beyond the scope of this paper. Nor do I intend to explain here *why* states would have an interest in reaching nonbinding agreements. My arguments are focused on trying to articulate the kind of practical reasons that such nonbinding agreements give rise to, and the latter would not normally depend on the particular reasons for reaching the agreement to begin with.

However, there is one important exception. Agreements, international treaties included, sometimes aim to solve a recurrent coordination problem between the relevant parties. When there is a standard coordination problem to solve, the incentives to abide by the solution are present ex-ante; a coordination problem occurs when both parties would rather act in concert with the other, and what they need for achieving this is reliable information about the other party's choice of conduct. The agreement, in such cases, simply provides the requisite information. Therefore, what matters for the solution of a standard coordination problem is only a reliable signal, that is, information about the parties' choice between the relevant alternatives. Since both parties have a dominant preference to act in concert with the other, once they have the reliable information about the choice between the alternatives, it doesn't matter whether the choice agreed upon is undertaken as an obligation or not. For the purposes of solving a coordination problem there is no need to undertake an obligation; the point of the agreement is to provide information.²⁸ In other words, when the reason for parties to sign an international treaty (or part of it) consists in the need to solve a recurrent coordination problem that arises between them, the agreement might as well be nonbinding, it makes no practical difference. To what extent, practically speaking, the coordinative rationale applies in the domain of international treaties I would not claim to know. But my point is that when the rationale applies to horizontal soft law, it provides a good explanation of the nonbinding

²⁸ Reality is much more complicated of course. My discussion assumes that the coordination problem is of the *pure* kind, where parties are largely indifferent between the feasible alternatives of conduct available and they would just need to know which one of these alternatives is chosen by the other. However, many coordination problems are not *pure* (modeled on the so called "battle of the sexes" game type situations). In such cases, it is still true that the dominant preference of the parties is to act in concert with the other, but they are not indifferent with respect to the choice between the feasible alternatives. In such cases reaching an agreement might be more difficult, and of course it opens the door for strong parties to exert pressure on the weak to follow their choice of the solution to the coordination problem. But even if there is some pressure involved in the negotiating stage, once an agreement is reached and it provides the requisite information, both parties would have a dominant preference to cooperate.

character of the legal instrument. Therefore, the rest of my discussion of horizontal soft law assumes that the nonbinding agreement between the states is not there to solve a standard coordination problem. Those are sufficiently accounted for, in terms of the practical reasoning involved, by the coordinative rationale of the agreement.

Back to our initial question, then. What is the point of a nonbinding agreement (of a non-coordinative type) between A and B to φ in C? Each party commits to the other, I suggest, to regard its reason to φ in C as a presumptive reason. Namely, they both acknowledge *vis a vis* the other, and thus render salient amongst them, the fact that they have a reason to φ in C. They also acknowledge that the reason to φ is defeasible, but they undertake the commitment to explain to each other if and why the reason to φ happens to be defeated. And thus they open themselves to warranted criticism, perhaps even some pressure, in case they fail to φ or else, fail to demonstrate that φ -ing is truly defeated under the circumstances.

Presumptive reasons for action, as we saw earlier, are much weaker than obligations. To acknowledge that I have a presumptive reason to φ is quite consistent with an explicit repudiation of my obligation to φ . But notice that when presumptive reasons are incorporated into an agreement or a contract, they acquire an additional directional character, making it the business of the other party to demand either compliance or demonstration that the alleged defeating reasons are convincing. This is what I mean by the accountability relations that nonbinding agreements establish.

6.

Normally, people do not need to answer to others when they fail to comply with a reason for action that applies to them. Reproaching someone for failure to comply with a reason for action, whether obligatory or not, has to be done for a good reason, it is not something we are permitted to do without some standing to do it. “Mind your own business is” is the appropriate reaction to reproach without standing. But the standing to have a say about another’s compliance with reasons is not confined to obligations. Obviously, a directional obligation, owed by A to B would normally give B the standing to demand compliance or at least to reproach A for non-compliance. But this is not a necessary condition. People are sometimes permitted to reproach others for failure to

comply with a reason for action when something makes it their business to do so. And there are many reasons that can give people such standing. For example, parents would normally have a standing to demand compliance or, at least, express some reproach for non compliance, with all sorts of reasons for action that apply to their children, and they have this standing in virtue of their parenthood. Sometimes special circumstances give people the standing to reproach others for failure to comply with reasons. Consider, for example, someone who happens to be in a unique position to help you in extreme circumstances of a dire need and fails to do so for no good reason. You may have the standing to reproach them. Either way, expressing demands on another's compliance with reasons for action that apply to them has to be done for a good reason.²⁹

Authorities, as we noted earlier, would have a standing to demand compliance with their legitimate directives even if the obligation they purport to impose is not construed as a directional one, owed to the authority as such. And the same applies to authoritative advice. As I argued earlier, the warranted expectation of a subject to demonstrate that the first order reasons are defeated, when they are, is precisely what makes authoritative advice different from any other advice; it constitutes a justified expectation of either compliance or demonstration that reasons to comply are defeated in the particular circumstances. In other words, a legitimate authoritative advice makes it the business of the authority to have a say about compliance. Something very similar is created by a nonbinding agreement. The agreement makes it the case that each party can legitimately expect the other party either to comply or to demonstrate that their defeating reasons for not complying are compelling. When there is no authoritative relationship between the parties, they can constitute a similar relationship contractually, as it were, by mutually rendering themselves accountable to each other. And this accountability relationship is, paradigmatically, directional, owed by one party to the other and *visa versa*.

A clarification is called for at this point. Earlier I expressed some doubts about the idea that reasons (or duties) for complying with authoritative directives are generally directional, owed by the subject to the relevant authority. But agreements are crucially

²⁹ In this I am expressing agreement with Raz's critique of Darwall, see Raz, "On Respect, Authority and Neutrality", at 295.

different. It is probably the case that any agreement (or contract) necessarily establishes mutual accountability between the parties to the agreement. Before you and I reach an agreement it may not have been any of my business whether you φ or not; but if we make an agreement that you need to φ , even if the φ -ing is not for my benefit, we thereby make it my business to have a say about it. If the agreement is meant to be binding, it would normally give me a right to your compliance and, typically, some compensation for failure to comply. If the agreement is nonbinding, it would give me the standing to have a say about your compliance, typically by making it appropriate for me to demand either compliance or demonstration of reasons allegedly defeating compliance.

Generally speaking, the level and content of the accountability relations created by agreements vary with different types of agreements, depending on their binding or nonbinding character and other content dependent factors.³⁰ But some level of mutual accountability is what agreements *inter alia* constitute. At the minimum, in case of nonbinding agreements, mutual accountability makes it appropriate for each party to express a direct interest in the other's compliance, typically a warranted expectation to get the relevant information about it, and in case of failure to comply, to demand an explanation, seeking a demonstration that there are indeed good reasons defeating compliance.

You might think that my suggestion to see the reasons constituted by nonbinding agreements as mutual recognition of presumptive reasons for action is too weak. So perhaps it might be helpful to dwell on the practical ramifications of such weak agreements, seeing that they are not necessarily negligible. To begin with, creating the standing to express demands in such contexts, is not without practical significance. When one is justified in expressing demands that the other comply or else demonstrate that non-compliance is supported by compelling defeaters, one would be justified in employing various reasonable means to try to make such demands effective, if need arises. Reasonable means would certainly include, for example, the appropriateness of seeking to get relevant information from the other party. But it may justify more than that. When A fails to comply with the presumptive reasons to φ in C, B might be justified in shaming

³⁰ Bear in mind that some contracts create obligations towards third parties, and many contracts allow for the transfer of contractual rights to third parties. In other words, the directionality of duties created by binding agreements are not exclusively bilateral, they can be much more complex.

A in public, maybe even withholding related benefits (even unrelated?) from B if A is in a position to do so, and in some cases, perhaps even exerting direct pressure to comply.³¹ What counts as reasonable means of making demands in these contexts is difficult to generalize. I would assume that the more voluntary and fair the initial agreement in question is, and the more is at stake, the more pressure one party can exert on the other to comply or else demonstrate that compliance is indeed defeated by other reasons.

Another practical ramification of nonbinding agreements pertains to mechanisms of monitoring compliance, often accompanied by ex ante agreed upon standards or measurements for what counts as adequate compliance. This is particularly true about soft law in the international domain. Treaties and authoritative nonbinding directives often include an elaborate set of standards for compliance, and various monitoring devices, that parties can employ to make their demands of each other explicit, less disputable, generally more transparent. I'm not suggesting that such mechanisms are necessarily or even typically effective. But they serve the dual purpose of reaffirming the issues of standing, and of rendering the monitoring of compliance more structured and publically accessible.

Finally, a practical import of nonbinding agreements, even if they do no more than creating presumptive reasons, consists in the moral vindication of complying parties. It enables those who comply with the agreement to claim the high moral ground, so to speak. Whether this has any real beneficial effects in the international domain is questionable, of course. One would have to assume that there are huge variations in the circumstances bearing on the relative importance of such issues.

Let me conclude by qualifying the moral of the story. My aim in this essay has been limited to showing that nonbinding legal instruments, of the kind we call soft law in the international domain, can be explained in terms of the practical reasons such instruments constitute. But to give a practical reason account of a practice does not mean that you commend it, morally or otherwise. It has not been my intention to implicate that the proliferation of soft law in the international domain is generally a good thing (or bad for that matter). Trying to answer that question would require taking into account a very

³¹ States and other international actors, including NGOs, do not shy away from pressuring states to comply with international treaties even if the latter are nonbinding. Sanctions are often employed, though not necessarily in a fair and equitable manner, needless to say.

wide range of considerations, many of them empirical, others normative, and some very difficult to ascertain. A tall order this is, requiring a different methodology from the one I used here. I can only guess that the answers would not be uniform across the board.³²

³² I am very grateful to David Enoch, Ori Herstein, Jens Ohlin, and the participants of the Hebrew University Law Faculty Workshop for helpful comments on a previous draft.