Discussion

Laura Valentini*

Response to My Critics

https://doi.org/10.1515/auk-2024-2023

Abstract: In *Morality and Socially Constructed Norms*, I argue that norms that exist as a matter of social fact have moral force, when they do, by virtue of what I call the 'agency-respect principle.' In what follows, I address the comments and criticisms of my view kindly offered by N. P. Adams, Åsa Burman, George Klosko, Katharina Nieswandt, and Titus Stahl, and which have appeared in a previous issue of this journal. My responses, just like the corresponding criticisms, will address some of the core themes of the book, including: the nature of socially constructed norms, the plausibility of the agency-respect principle, how to best understand and ground political obligation, and, in general, whether we should think that socially constructed norms have (primarily) moral normativity, as opposed to some other type of normativity.

Keywords: socially constructed norms; agency respect; political obligation; political realism; moral normativity; joint commitments

Let me start by sincerely thanking the colleagues who have contributed to this symposium for their insights and generosity. I have learnt a lot from their critical engagement with my work. In this response, I will proceed as follows. I will first address Katharina Nieswandt's comments, which challenge the very plausibility of the principle lying at the heart of my view: the agency-respect principle. I will then turn to Åsa Burman's and Titus Stahl's discussions, which raise doubts about the way I characterize socially constructed norms. Next, I will engage with George Klosko's piece, which focuses on a particular application of the agency-respect view, namely that to political obligation. I will then conclude by discussing N. P. Adams' article, which offers an external critique of the position I defend, questioning not so much the details of my view, but rather the broader philosophical perspective from which it is developed. In order to avoid repetitions, in what follows, I assume that the main lines of argument of the book are known to readers. For a summary of the

***Corresponding author: Laura Valentini**, Fakultät für Philosophie, Wissenschaftstheorie und Religionswissenschaft, LMU Munich, Munich, Germany, E-mail: l.valentini@lmu.de

book's contents, I refer readers to the précis published in this journal's previous issue (Valentini 2024).

1 Is the Agency-Respect Principle Empty or Implausible? A Reply to Nieswandt

On my view, the moral normativity of socially constructed norms is derived from the agency-respect principle (ARP). According to this principle, we have a pro tanto obligation to respect people's authentic commitments, provided these are permissible, and doing so is not too costly to ourselves. Since, I argue, socially constructed norms are constituted by people's commitments, the agency-respect principle also requires us to go along with those norms, provided the applicable disclaimers (authenticity, permissibility, etc.) are satisfied.

Nieswandt's (2024) challenging discussion targets precisely this principle, hence the core of my view. She expresses general skepticism about the Kantian tradition where my view is situated, and finds that the agency-respect principle *both* lacks intuitive plausibility *and* is, ultimately, empty (Nieswandt 2024, 53).[1]

To support her worry about *implausibility*, Nieswandt invites the reader to consider the following scenario:

> Suppose my neighbour has the quirky, idiosyncratic habit of parading our street every morning dressed in nothing but the national flag (or make that the pride flag if you prefer) and would appreciate it if passers-by briefly stood to attention. Suppose furthermore that the patriotic (or the queer) cause is permissible and that leaflets in our mailboxes have made all residents aware of our neighbour's desire. (Nieswandt 2024, 54)

Nieswandt then worries that, according to the ARP, passers-by who are aware of the neighbour's desire have a pro tanto obligation to stand to attention, provided this is not too costly to them and that the neighbour's desire is permissible. Now, I agree

1 Nieswandt's reconstruction of my view contains a couple of imprecisions. Even if these do not affect her main lines of criticism, I nonetheless briefly address them in this footnote. First, Nieswandt often says that the ARP gives us duties to respect people's convictions (and at times she speaks of desires or preferences). This is not quite right – the ARP requires us to respect people's *commitments*, namely their robust intentions to pursue particular ends, goals, values, etc. Second, Nieswandt (2024, 55) attributes to me the view that the ARP generates obligations only when commitments are constitutive of socially constructed norms, and rightly wonders where this scope restriction comes from. This is a misunderstanding – the very formulation of the ARP shows that it applies whenever commitments of the right sort are in play, whether constitutive of socially constructed norms or not. The subject of the book is, of course, the moral normativity of socially constructed norms, but the principle itself is meant to be fully general, and its application not restricted to that domain.

with Nieswandt that this conclusion is counter-intuitive, but I do not think it necessarily follows from the ARP. First, Nieswandt talks about desires, while the ARP focuses on commitments, which involve a much stronger agential involvement than desires – indeed, in my characterization of socially constructed norms in the book, I explicitly exclude desires as part-constituent of them (Valentini 2023, 33ff.). Second, even if we switched to the neighbour's commitments, considerations of authenticity and costs would likely come into play. Given how idiosyncratic this neighbour's commitment would be, can we genuinely speak of authentic agency here? Or is the neighbour actually disturbed, and acting out of some compulsion? And would it not be likely too agency-costly for passers-by to stand to attention, especially if they are not patriotic themselves? In sum, I am not sure the neighbour's example is ideally suited to prove the point Nieswandt wishes to prove. But I still think she is onto something, namely the possibility that, following the ARP, we may have pro tanto obligations to respect people's commitments in cases where this seems rather odd. I think she is right about this, but I am less troubled by this conclusion than she is. I explain why in what follows.

We find it (I assume) not implausible that, even if we are not religious believers, there would be something disrespectful in, say, entering a temple wearing nothing other than a bikini and a baseball cap. The victims of disrespect, in turn, would be the relevant religious believers, not the temple's deities (whose existence an outsider may not believe in). It is the believers' commitments we are disrespecting, hence their agency and the normative world they have built. In the familiar case of religion, the agency-respect principle seems to do well: its implications do not appear as implausible as in the neighbour's case. But suppose now we look at an individual's somewhat eccentric commitment, say, their commitment to spending multiple hours every day observing butterflies in nature. It does indeed follow from the ARP that, provided this commitment is authentic, permissible, and respecting it is not too costly, those around our butterfly-lover should respect it as well as they can. This may mean not interfering with it, and possibly also making adjustments to accommodate it. Perhaps puzzling at first, I think this follows from what taking others seriously as agents implies. We look at them as shaping their lives and giving meaning to them through their commitments, individually and together. Even if we do not share those commitments, and even if we do not find them valuable ourselves, we see the value *in the person* whose commitments they are and understand that those commitments are an expression of their agency, and so worthy of respect *for that reason*.

This line of reasoning is premised on a certain conception of the person as autonomous agent, worthy of respect as such. I treat this, as Nieswandt notes, as an 'axiom' in my account. Of course, if one finds this Kant-inspired, person/autonomy/agency-centred perspective altogether unappealing, one will not

feel very moved by the view defended in the book. But one needs to start from somewhere in one's theorizing, and this seemed to me like a good place to start. In sum, I do not think the implications of the ARP are as implausible as Nieswandt makes them out to be, but I agree that the ARP might well require us to respect commitments we ourselves find of little (or no) value. This is because the genuine object of respect is not the content of the commitments, but the agent whose commitments they are.

Now let me move on to Nieswandt's concern about *emptiness*. This concern, as Nieswandt notes, is explicitly addressed in the book, so I will devote less space to it (Valentini 2023, sec. 8.2). Here the worry is that, since I leave the parameters of the ARP – namely authenticity, permissibility, and costliness – unspecified, the agency-respect view lacks any action-guidance capacity. Depending on how the parameters are specified, we will reach different conclusions: plug in different views about moral permissibility, or about what makes a commitment authentic, and the ARP will 'spit out' different verdicts.

Nieswandt's (2024, 56–9) analysis of how the ARP works is spot on, but I do not think this shows that the ARP is ultimately empty. For the ARP still captures an important truth, namely that respect for authentic agency matters, subject to permissibility and costs constraints. But to the extent that people (reasonably) disagree about permissibility and costs, different ways of specifying these parameters will satisfy different audiences. Is this a lethal blow to the ARP? I believe not. For one, there will be clear-cut cases in which people will converge on authenticity, permissibility, and costs. Here, the principle will deliver univocal answers. Then there will be trickier, marginal cases. But it cannot be the job of the ARP to adjudicate between them. For example, people reasonably disagree about whether it is morally permissible to eat non-human animals. The ARP, as formulated, is agnostic on this matter. So, say, a vegetarian who believes that eating animals means supporting mass murder will not consider a religiously motivated commitment to slaughtering and roasting a goat on a particular festivity permissible, hence a source of agency-respect claims. By contrast, a carnivore would likely have a different view. The ARP does not adjudicate the question of the moral status and appropriate treatment of non-human animals (among many other such questions), but I do not think it has to, in order not to count as empty.

Consider the following analogy. On some views, promises are binding in that they amount to our exercise of a normative power, but their bindingness is conditional on the promise's content being morally permissible. Are these views empty just because they do not come equipped with a full account of moral permissibility? Surely, the claim that the source of the moral bindingness of promises is the exercise of a normative power (as opposed to, say, an independent principle to honor legitimate expectations) has explanatory merits, even if it does not, per se, tell us which

promises are binding and which ones are not, unless coupled with a full account of permissibility (e.g., Owens 2006; Scanlon 1990).

In sum, while Nieswandt is right that the ARP might both deliver counter-intuitive results and give less guidance than one might have hoped for, I do not think either concern is sufficiently serious to undermine it.

I will now also comment, more briefly, on Nieswandt's (2024, 59ff.) critique of the ad-hocness charge I raise against her conventionalist view. Here, I must confess, I found myself a little lost. Her discussion, and my own reaction to it when I read it, suggest that we both feel that we have respectively misunderstood one another's arguments. I am sorry about this. But since I did not quite recognize my own critique of her view in what she has written as part of this symposium, I will give a quick summary of it, to try to see whether possible good-faith misunderstandings can be cleared. Of course, there's always the risk that I'll make things worse. (I hope not!)

In the relevant portion of the book where I discuss her view, I specifically focus on one of her articles about conventionalism, which I find in many ways illuminating. There, Nieswandt (2018, 17) says that, according to conventionalism: "If the rules of [a] practice … say that X must φ given fact F, and [the practice is justified], then, given F, X must φ because of F." Looked at from the perspective of the book's topic, this principle can ground the moral normativity of socially constructed norms – it says that, provided those norms (and the practices they constitute) are justified, one should do what the norms require, because they so require. The conventionalist principle enjoins us to take the practice at normative face value, provided it is independently justified.

In the 2018 piece, Nieswandt defends the conventionalist position against non-conventionalist alternatives on account of its ability to avoid the well-known 'wrong reasons problem' (Scanlon 1990). To see this, take the practice of promising. Some, for instance, argue that our duties to keep promises stem from our obligations to do our part in practices that are mutually beneficial, with promising being one such practice. The trouble is that this explanation does not really fit the practice – it suggests that by breaching a promise, we wrong every participant in the practice, and not the promisee in particular. A failure to 'do our part' seems to be the wrong kind of reason. Conventionalism avoids this difficulty. It says that, provided the practice of promising is justified (say, because it is mutually beneficial), *then* the reason why one ought to do what one promised is that one promised (thereby generating a duty for oneself, owed to the promisee).

The gist of my critique of the conventionalist view in the book is that it is *ad hoc* because seemingly reverse-engineered to respond to the wrong-reasons problem. Go back again to the case of promising. It is true that, if we treat the overall justification of a practice as a background condition that is part of the conventionalist principle, and not as the ultimate ground of practice-internal duties, we avoid the

wrong reasons problem. But the question is what reason do we have to do it, other than this allowing us to avoid the wrong reasons problem? *In the particular paper* by Nieswandt I discuss, I did not find a convincing answer to this question, specifically not one that would satisfy the 'deflationist,' who is the main opponent of those who think that socially constructed norms have moral normativity.

Put differently, the deflationist says (a) 'treat the considerations you invoke to justify a practice as the grounds of the practice-internal duties,' while the conventionalist says (b) 'treat those considerations as background conditions for the practice to be able to generate duties.' What the deflationist proposes, namely (a), seems prima facie very plausible; only it sometimes involves somewhat revisionist implications with respect to how we conceive of our duties (such as the wrong reasons problem). By contrast, (b) lacks this independent plausibility and seems constructed with the specific aim of avoiding the wrong reasons problem. This is the gist of my complaint. Here, I do not wish to defend the complaint – maybe I am wrong! My aim is just to clarify what the complaint is, and where the charge of ad-hocness for conventionalism comes from.

Importantly, Nieswandt has other work on conventionalism, which she discusses at some length in her contribution to this symposium, and *nothing* in what I say in the book is meant to target *that other work*. Maybe this is why I was puzzled by her reaction. My discussion targets the arguments in her 2018 article specifically, which I found spoke directly to the particular issues I was tackling in the book. Perhaps one could argue that my selection of work to rely on was not the best (I do not know), but it is important to bear in mind what the precise target of my discussion was. I cannot exclude that, had I taken a more comprehensive perspective on Nieswandt's conventionalism, I would have argued differently.

2 Are There Teleological Norms Proper? A Reply to Burman

Burman's (2024) discussion of my book challenges my characterization of socially constructed norms on account of it being too narrow and leaving out an important class of such norms. Burman suggests that there can be two types of socially constructed norms: deontic and teleological. Deontic norms set out standards of behaviour in the form of rights and obligations. Teleological norms, by contrast, set forth standards of excellence we should aspire to. Burman then goes on to argue that my book only really concerns itself with socially constructed norms of the first kind, and that expanding the account to teleological norms would require significant departures from the agency-respect view, more in the direction of virtue ethics.

To make Burman's point more concrete, consider the following two propositions. (One can tell I was in the middle of grading when writing this response!)

1. Professors ought to return grades to students within 3 weeks of submission, with at least one paragraph of feedback.
2. A good/excellent professor returns grades to students promptly, writes extensive feedback, and offers to discuss this feedback with them during office hours.

The first proposition is a deontic norm, the second sets out what Burman calls a teleological norm. Burman suggests that my view only accommodates 1, while in fact it should also accommodate 2. The reason, she points out, is internal to the view itself. This is because my interest in the moral normativity of socially constructed norms partly stems from the fact that violations of those norms trigger moral reactive attitudes such as resentment, guilt, and blame. And precisely these reactive attitudes, says Burman, can also be activated when the standards of teleological norms are not met (Burman 2024, 31). The question I raise in the book, therefore, arises for those norms too, and yet the book has very little to say about them.

I find Burman's comments interesting and helpful. In fact, they have given me a lot to think about. My immediate reaction to them is that I find the category of 'teleological norms' somewhat unstable. In particular, this category can be interpreted in two ways. On one interpretation, teleological norms collapse into deontic ones, hence my account straightforwardly applies to them too. On the other interpretation, teleological norms are not norms proper, in the sense I am interested in, hence violations of them do not trigger moral-style reactive attitudes. If so, then there's no reason, internal to my account, why they should fall within the scope of my analysis. (Which is not to say that they are not worth analyzing independently!)

Perhaps the easiest way to make my point is to start with a terminological observation. I find the contrast between deontic and teleological norms a little confusing, if used to mark the difference Burman is interested in. This is because deontology and teleology are both approaches to obligatory action. Roughly put, the former says that what we ought to do depends on the quality of the actions themselves, independently of consequences. The latter says that the normative status of an action depends on its consequences. So, to mark the distinction Burman wishes to draw attention to, I would prefer to employ the labels *deontic* and *evaluative*: the former having to do with what we ought to do, the latter with what is good, bad, excellent, etc.

Now, return to the case of professors. I have deliberately stated proposition 2 in non-deontic language, simply outlining what it takes for someone to be an excellent professor. It is a separate question, therefore, whether one ought to be an excellent professor. My thought is that if 2 is coupled with a background requirement that one ought to be an excellent professor, then 2 becomes a deontic norm, just like

1, violations of which trigger the relevant reactive attitudes. However, if there is no independent requirement to be an excellent professor, then falling short of the standard set out in 2 cannot elicit the same types of attitudes.

Moreover, I venture to hypothesize that the reason why we sometimes feel that falling short of standards of goodness or excellence merits reactive responses typical of deontic norms is precisely that, in the contexts in question, those standards have become the objects of deontic norms. In other words, if certain reactive attitudes are in place – namely guilt, resentment, and blame – this is a sign that a standard is not purely evaluative but has acquired a deontic character.

Let me illustrate the point more concretely. Consider two universities: Universities A and B (apologies for the lack of imagination). In University A, most professors return essays within a week, with 3–4 paragraphs of feedback, and offer plenty of office hours. A few professors, however, take 3 weeks to return the essays and only write one paragraph. Students receiving this more concise feedback feel resentful towards the slower professors, and their colleagues frown upon their behaviour, which they consider lazy and inconsiderate. In university B, by contrast, only a handful of professors return essays within a week of submission, and with very extensive feedback. Most of them use the full three weeks and write about one paragraph. The extra-generous professors are looked at with admiration (and, in some cases, disbelief) by their colleagues, and students typically express special gratitude to them, for the exceptional amount of time and effort they put into their grading. Yet colleagues who stick to the 3-weeks/one-paragraph policy do not feel like they are falling short or letting anybody down.

Faced with these two cases, I would be inclined to say that, in University A, there is a deontic norm requiring professors to be 'excellent.' By contrast, in University B, there is an evaluative standard (not a norm) for professors, satisfying which is considered generally praiseworthy, but not obligatory.

It may be that, in the real world, the line between deontic and evaluative is sometimes blurred, such that we cannot easily classify standards as falling under one or the other. But there seems to me to be a clear difference between the two. If I am right about this, then there are no reasons, internal to my account, why evaluative standards should also fall under the purview of the agency-respect view. This is because falling short of those standards does not elicit typical 'moral' reactive attitudes – such as guilt, resentment, and blame. It does of course elicit emotional responses, but these are responses of a different kind, perhaps more akin to those we face in cases of supererogation. And when the responses resemble those apt in cases of moral violations, then we have reason to believe that the evaluative standards in question are no longer purely evaluative but are in fact acquiring a deontic character. This, at any rate, is how I would respond to Burman's very well-put

challenge. No doubt, however, there is more to say. And possibly, further reflection might bring me to revise my own response!

3 Reviving the Joint-Commitment Account? A Reply to Stahl

Stahl's (2024) critical discussion of my view takes issue with its individualistic nature. Stahl rightly notes that, on the agency-respect view, the normative status of individual commitments and joint ones is identical. Following the agency respect principle (ARP), the reasons we have for respecting individual commitments and the commitments underpinning socially constructed norms are identical. Indeed, I claim that socially constructed norms are precisely made up of individual normative beliefs and commitments, which enough people in a given context hold as a matter of common knowledge. Stahl skillfully argues that this way of looking at socially constructed norms is excessively individualistic. Unless we recognize that the commitments making up socially constructed norms are *joint* in a sense stronger than being simply shared by enough people as a matter of common knowledge, we cannot account for a number of plausible intuitions about socially constructed norms and their moral normativity. Here are some examples Stahl offers, which are meant to illustrate this point. (There are more in his paper, but due to space constraints, I cannot address all of them here.)

Stahl asks us to imagine a university (UA) where participants in a seminar each have an individual belief and commitment to speakers wearing ties, and this is publicly known among them. He also asks us to imagine a similar scenario, taking place at a different university (UB) in which the host informs the speaker that "there is a *shared* and *jointly accepted* social norm at the university that prescribes wearing a tie" (Stahl 2024, 73, emphasis original). Stahl's contention is that an eccentric and anti-conformist academic invited to give a paper at these institutions would have a stronger reason to respect the local norm in the latter case than in the former, and this is because of the more properly joint nature of the commitment behind the norm.

Similarly, Stahl contends that we can distinguish between publicly accepted overlapping commitments to norms and properly joint such commitments by looking at the types of accountability structures that they respectively give rise to. So, suppose that I am a visitor at UB, and that in my home university (UC) the same norms as those jointly accepted at UB concerning wearing ties at seminars apply. Still, it would seem that members of UB would have a special standing to criticize one another for violating the relevant norm, which I lack. In other words, members of UB are united by normative bonds which do not encompass me, something

that the agency-respect view cannot capture, given that I happen to have the same commitments as them (with regard to wearing ties at seminars) and this is publicly known (Stahl 2024, 76).

Relatedly, suppose that the fact that UB and UC both accept the tie norm is common knowledge across members of both universities. It follows from my account of socially constructed norms, Stahl says, that the 'we' sustaining the norm in question spans across both universities (Stahl 2024, 76–7). But this seems counter-intuitive, for in fact we would want to say that there are two 'we-s,' which each happen to accept the same norm, and this is mutually known among them: members of UB and members of UC.

Stahl suggests that, in order to avoid these difficulties, my view of socially constructed norms needs to rely on a thicker account of joint commitments. Yet he agrees with me that the prominent account of joint commitments offered by Gilbert (2006) would be too demanding. In a nutshell, Gilbert understands those commitments on an agreement-based, quasi-contractual model, which – as I and others have remarked – cannot plausibly be scaled up to the types of social norms that govern large-scale communities. Stahl then creatively and constructively offers an alternative account of 'jointness' which, he says, can be used to enrich my characterization of socially constructed norms, thus allowing it to avoid the difficulties previously highlighted.

My original account says that, for a socially constructed norm that one ought to φ in a given context C to exist, the following conditions have to obtain:

- *Belief condition*: enough people in C believe that one ought to φ.
- *Commitment condition*: enough people in C are committed to 'one ought to φ' functioning as a general standard of behaviour.
- *Publicity condition*: the fact that the belief and commitment conditions are satisfied is common knowledge among members of C.

Stahl's proposed amendment retains the belief and commitment conditions but replaces the publicity condition with what could be named a 'mutual accountability' condition.

- *Mutual accountability condition*: enough people in C understand themselves to be liable to criticism from each other and to be entitled to criticize each other for failures to comply with 'one ought to φ' qua 'representatives of the group' (Stahl 2024, 79).[2]

2 Stahl talks about a group G, rather than about individuals in a context C, but this is a rather significant departure from my own way of looking at things, and one that he does not acknowledge in his discussion. The point about mutual accountability he does want to make, however, is preserved if one refers to C rather than a group G.

With this mutual accountability condition, says Stahl, we are now equipped to respond to the challenges he considered earlier. We can explain why it still makes sense to say that the 'we-s' supporting the tie-norm at UB and UC are different. We can make sense of the greater strength of the non-conformist academic's obligation to wear a tie in UB rather than in UA. This is because the joint commitments underpinning UB – ones that come with special accountability relations – matter greatly to the agency of participants, since their own identity is shaped by their belonging to communities sustained by practice-internal obligations. And similarly, we can make sense of the greater strength of members' obligations towards each other since, as co-members, they need to sustain communities that give access to practice-internal reasons central to their identity (Stahl 2024, 82).

Stahl's comments are excellent and have given me a lot to think about. In addressing them, my strategy will be twofold. I will try, on the one hand, to blunt the intuitive force of some of his objections and, on the other hand, to show that his proposed solution only works assuming the strong, Gilbert-style account of joint commitments he himself finds implausibly demanding.

Let me start by going back to Stahl's case of an academic arriving at a university where the host announces that "here we have a norm requiring speakers to wear ties." Now, unless this is interpreted as "here we have *agreed* to adopt the norm that speakers should wear ties" (hence, the norm originated in a quasi-contractual fashion), it is not clear how the statement should differ from the claim that "this is what we do around here, and each knows that the others know that this is what we do around here, and so on." The host's statement is intuitively accompanied by greater normative force if it is implicitly underpinned by the contractual model.

Second, Stahl remarks that members of the same group stand in stricter mutual-accountability relations when it comes to 'their own norms.' But why this should be so, in the absence of a contractual-like joint commitment, remains opaque. To see this, consider again the queuing norm. Suppose I am a foreigner visiting the UK, and I independently happen to subscribe to the queuing norm. When I see someone skipping the queue, I reproach them and display the kinds of accountability responses that British people typically exhibit. On Stahl's view, however, British people have greater standing to hold queue-jumpers accountable in a case like this than me, since they (but not I) belong to the 'we' supporting the norm being violated. Yet, in the absence of a quasi-contractual transaction among them, one in which the British have agreed with each other to support the queuing norm, it is unclear where this greater accountability is supposed to come from. If I support the norm as much as they do, but none of us has engaged in any contract-style transaction, then why am I excluded from the relevant 'we'? I understand that intuitions suggest that I am, but perhaps they should not. Or rather, I rightly

do not belong to the 'we' of the 'British people,' but I do belong to the 'we' of the 'queuing-norm supporters.'

Finally, I should emphasize that I am not denying that membership may – even short of contract-style mechanisms – generate obligations. In the book, I remain agnostic on this matter – and here, of course, a lot depends on what exactly makes one a member. But the question then is whether the obligations in question are of the kind I am trying to establish – namely content-independent obligations to do as the norms require, because they so require, and not obligations that ultimately stem from the merits of the norms under analysis. As I explain in the book, unless the norms stem from contractual-style mechanisms, it is unclear that alternative grounds for obligation really do deliver the desired content-independence. This nicely leads me to George Klosko's comments. But before addressing them, I should emphasize again that I have found Stahl's discussion particularly useful and challenging, and that what I have offered here is only a preliminary and certainly still incomplete response to it.

4 A Sophisticated Deflationary View? A Reply to Klosko

Klosko's (2024) discussion focuses on a particular application of the agency-respect view, namely that to political obligation. On the agency-respect view, individuals' obligations to obey the law 'because it is the law' stem from agency-respect vis-à-vis citizens' commitments to the rule of law. On this view, while there often are many content-dependent reasons why we should do what the law requires – e.g., what the law mandates is instrumentally useful or already prescribed by morality – there is also a content-independent such reason (in fact, a pro tanto obligation), namely that abiding by the law is necessary to respect citizens' commitment to the rule of law. This pro tanto obligation is content-independent because what grounds it is not the substantive goodness of the law, but the fact that legal requirements, qua legal requirements, are the object of certain kinds of commitments.

Klosko argues that while my account – unlike other prominent accounts in the literature – successfully delivers a content-independent obligation to obey the law, it is less persuasive than what I call 'deflationary views,' namely views which ultimately trace the moral force of law to features of its content. What is more, Klosko suggests that the superiority of deflationary alternatives can be established relative to the criteria I myself employ to demonstrate the plausibility of my own agency-respect view, namely explanatory power and fit. For Klosko, there exist versions of the deflationary view that do better than the agency-respect view on both counts (Klosko 2024, 43ff.). These are versions of the deflationary view that acknowledge

uncertainty and disagreement about what the law should be, and emphasize that part of what makes the law instrumentally useful is its providing cognitive shortcuts and coordination points for its subjects.[3] So, for instance, whenever I am faced with a red light, or with a tax bill, or with questions concerning a floor's minimum carrying capacity, I do not need to independently work out whether driving through would be efficient and safe, whether what I am paying is fair, or what safety margins are best overall. In these various domains, it makes most sense to defer to the law, provided the law does a good enough job at tracking relevant applicable considerations (Klosko 2024, 45).

This version of the deflationary view, says Klosko, works better than the agency-respect view in terms of *explanatory power*, since, when it comes to the phenomenology behind people's sense of duty to the law, this often refers back to the law's content, to what it does for us. We do not pay taxes primarily out of respect for others' commitments to the rule of law, but because we feel this promotes important social goals. We do not stop at red lights because this is what respecting our fellow-citizens' commitment to the rule of law demands, but because it typically makes good prudential sense (Klosko 2024, 41–2).

Moreover, the view also does well in terms of *fit* with our considered judgements about when we are morally obligated to abide by the law. This is because, Klosko notes, on my view we have a pro tanto obligation to do what the law requires 'because it so requires' only when this is not too costly to our own agency. However, many laws, from tax laws to forced conscription laws, are going to be quite costly. And yet, these are precisely the kinds of laws that a successful theory of political obligation should be able to demonstrate place obligations on us. The modified deflationary view is successful in this respect, since it appeals to weighty content-dependent reasons. But the agency-respect view, with its cost-constraint, seems to struggle to come to the same conclusion (Klosko 2024, 46–7).

In sum, Klosko's view is that we should be deflationists about political obligation after all, and that doing so still allows us to account for most of our judgements about when the law has moral force, and why – in fact, better than the agency-respect view does.

I must confess that I agree with many of Klosko's remarks. In fact, as he also notes, I explicitly say that the weightiest reasons to do as the law requires tend to be content-dependent, namely the types of reasons invoked by views located in the deflationary camp. The question, for me, is whether there also exist content-independent reasons, and if so, what these reasons might be. The agency-respect view is supposed to offer an answer to this question that supplements, rather than

3 This line of reasoning seems to me to be reminiscent of the rationale behind Raz's (1985) service conception of authority.

competes with, deflationary views. From my perspective, deflationary views have a lot going for them, but they are not the whole story. Therefore, for Klosko's critique to be genuinely threatening to my account, it would have to show that a deflationary approach explains *all that we wish to explain*. I believe it does not. Moreover, I think the 'fit' difficulties Klosko detects in the agency-respect view are not there.

First, even Klosko's sophisticated version of the deflationary view seems to me to struggle to fit our intuitions about cases where violating a norm or a law seems intuitively wrong, and yet it is obvious that doing so violates no moral principle and leads to no problematic consequences. Consider instances of harmless trespass, where it is clear that nobody will be harmed, offended, and the like. Why should we, in such cases, still defer to the law and not trespass? The epistemic or coordination benefits of doing so do not seem to provide a convincing explanation. Similarly, if we could secretly walk into a temple wearing forbidden footwear without anyone noticing, my sense is that we'd still be doing something wrong – but the explanation cannot quite rest on the benefits of going barefoot for, ex hypothesi, there are none. These strike me as cases in which the deflationary view, no matter how sophisticated, does not quite fit the evidence. From a deflationary perspective, we have good reason to trespass and to wear shoes in the temple, if the circumstances allow us to do so without violating any other prudential or moral principle. But if we think that such behaviour is morally problematic, then we need to find an explanation for it, and the agency-respect view provides just that. Given the generality of the agency-respect principle, this explanation in turn generalizes to all sorts of norms supported by genuine and permissible commitments.

Does this mean that, on the agency-respect view, the main reason we should, say, pay taxes and stop at red lights is agency-respect for others' commitments to the rule of law? No. As I readily admit, in many cases, agency-respect will only cover a small fraction of the normative force of law-mandated/prohibited conduct, with the bulk of it being accounted for by the kinds of reasons deflationists invoke. This is why I say that the agency-respect view is, after all, not too distant from philosophical anarchism. Perhaps this is disappointing, but it is the way it is.

Let me now turn to Klosko's worries about the agency-respect view's inability to ground obligations to do what the law requires when what the law requires is rather costly. Two brief remarks here. First, as already noted, the agency-respect view offers additional reasons to those provided by the deflationary view for why we should act as the law requires. So, even if the agency-respect view were not able to ground costly obligations, this would not imply that those obligations do not exist. They may well exist but have different grounds. Second, while I do not provide a criterion for what counts as too costly, in at least one of the examples mentioned by Klosko, it seems to me that a plausible specification of costs would

allow the agency-respect view to come to plausible conclusions. This is the example of taxation.

Provided the relevant tax system is sufficiently fair, then no individual contributing to it can plausibly claim that their contribution is too restrictive of their own agency, that they are making too big a sacrifice in terms of their agency, in order to support or accommodate that of others. Extortionate taxes could be the target of a suitable agency-cost complaint; fair taxes, I would think, not. Things are trickier with Klosko's other example, namely conscription. Here I would be inclined to say that the costs to one's agency may indeed be too high for one to have agency-respect obligations to enroll in the army, but this seems to me exactly right. The institution of conscientious objection arguably reflects this very fact. For a pacifist, serving in the army places too big a burden on their agency, which seems to be a rather good reason for making an obligation of this kind at the very least opt-out.

In sum, I think Klosko is right in arguing that a sophisticated version of the deflationary view can explain much when it comes to our obligations to do what the law requires, but I do not think it can explain all. The agency-respect view can fill this gap (and other gaps, when it comes to non-legal, but still socially constructed, norms).

5 A Flattening of Our Normative Landscape? A Reply to Adams

Adams (2024) offers an *external critique* of the position defended in the book, seeing it as a representative example of philosophical theorizing from within what Bernard Williams calls the 'morality system.' Adams finds the very framing of the question I raise – i.e., do we have *moral* obligations to do/refrain from doing certain things because socially constructed norms require us to do so? – problematic. It is problematic, Adams suggests, because it presupposes that all valid 'oughts' are moral ones, and thereby flattens our picture of normativity. For, in addition to moral normativity, there is normativity that is "ethical, prudential, characterological, and much else" (Adams 2024, 15). Perhaps, these other types of normativity are in play in the case of socially constructed norms. Why assume that moral normativity is what matters here? Doing so, Adams says, distorts the nature of the phenomena I am trying to investigate. Indeed, given how diverse socially constructed norms are, we should expect their sources of normativity to vary, and not be just a matter of morality. Adams makes these points in general, and then illustrates them more specifically with respect to legal norms.

Before reading Adams' informative discussion, I had not come across Williams' idea of a 'morality system.' Now, if being committed to the morality system means

denying that there exist forms of normativity other than moral normativity, then I have yet to meet someone who embraces such a commitment. Adams seems to infer from the fact that I ask whether socially constructed norms have moral normativity that I am committed to the view that moral normativity is the only normativity they could possibly have. This does not follow. Trivially, socially constructed norms have prudential normativity – one often has very good prudential reasons to conform – and nothing in the book, I think, suggests otherwise. I am simply interested in whether they also have moral normativity qua socially constructed norms, and not merely due to the merits of the actions they prescribe.

Now, one might ask why I am interested in this question specifically – and perhaps this is where the real disagreement between Adams and myself lies. A key driver of my interest is the fact that we often respond to violations of socially constructed norms – from the queuing norm to politeness norms all the way to legal norms – in ways that are typical markers of moral violations. That is, we resent violators, violators feel guilty, violators are seen as fitting objects of blame, etc. I assume that these types of reactive attitudes/emotions are warranted whenever we are in the presence of wrongdoing. So, to me, asking whether those reactions are warranted is asking whether there is a moral violation lying behind disobedience to socially constructed norms 'qua norms.'

Adams (2024, 17) challenges this line of thought, saying that: "Warranted feelings of guilt, shame, resentment, and so on can track social oughts and ethical oughts, not only moral oughts." And, in the subsequent paragraph, he remarks – again echoing Williams: "I think we can be warrantedly called to account within social practices without implicating our most basic claims on each other as free and equal persons, so without even raising the question of moral justification."

These two passages reveal that Adams and I circumscribe the domain of the 'moral' differently. I think of the domain of the moral as involving – roughly speaking – rules of conduct about how we should relate to each other, the violation of which warrants certain kinds of reactive attitudes. What considerations, in turn, make those reactive attitudes warranted, is a further matter. Adams, by contrast, seems to think that morality involves rules of conduct that implicate our most basic claims on each other as free and equal persons.

I do not think this is the place to adjudicate between two different characterizations of what falls within the domain of the moral. But I will at least try to motivate my own characterization. First, I wish to remain agnostic between different substantive accounts of morality (universalist, particularist, etc.). Ideally, I'd like to adopt a definition of the moral domain that can allow us to see how a utilitarian, a Kantian, a universalist, a particularist, etc. are all busy trying to figure out what morality requires, and not talking past each other. Second, I want to allow for the intelligibility of common-sense claims such as – 'violating one's promises is morally

wrong,' or 'letting one's friends down is morally wrong,' etc. – even if it is unlikely that either behaviour implicates our most basic claims as humans.

It is true that this leads to broadening the realm of the moral compared to a characterization just focused on our claims as free and equal persons. But it is not clear why this is a bad thing. Adams (2024, 16), for instance, asks "someone who jumps an English queue is certainly a jerk, but is he morally wronging others or is he … just a jerk?" Similarly, Adams (2024, 16, 18) remarks that we can legitimately criticize others for being impolite or for cheating, without implicating that any moral violation has taken place.

Maybe this is just me, but I do not quite see the problem with saying that a jerk behaves immorally, or that someone who cheats does something morally wrong. And, with politeness too, to the extent that violations of (non-pernicious) politeness norms convey disrespect, then I think there's nothing odd in concluding that such violations are morally wrong. To be sure, they are not as wrong as human rights violations! A jerk and a genocidal dictator are not morally on a par. But this is consistent with both of them behaving immorally, hence warranting certain kinds of reactive responses – though clearly of different gravity/intensity.

In short, I am not sure Adams and I really have a genuine substantive disagreement, rather than a purely semantic one, concerning what counts as part of the realm of morality. And as far as the latter is concerned, without being told more, I am not sure why Adams' Williams-inspired more restrictive (and substantively committed) account should be preferred to the one I assume in the book.

Perhaps the answer is to be found in Adams' critique of my take on political obligation and the authority of law. In the book, I argue that our pro tanto obligation to obey the law because it is the law is grounded in a broad moral principle, namely the agency-respect principle. Adams again offers an external critique of this way of proceeding, suggesting that approaching the question of political obligation by appeal to an external moral principle is virtually self-defeating. He does not specifically target the principle I offer, but the very framing of the question and approach to answering it I adopt. Once again, his point of reference is Williams' political realist approach.

I am not sure I'll be able to offer a satisfactory response to this portion of Adams' critique, partly because I feel that such a response would require me to explain why I am not a political realist à la Williams. Here, the honest answer is: I find political realism very intriguing, but I never really fully understood it and, more importantly, I never understood its purported ability to avoid 'moral' considerations in thinking about politics. I could go on a while substantiating these claims, but I am not sure this is the right place to do so (moreover, there are several critiques of the realist position that already make them).

I will thus conclude by addressing what, I think, is the main substantive realist-inspired worry Adams raises against my in fact, any non-realist's – position on political obligation. This is that it has no room for the reality of "separate persons forging political consensus by exercising private political judgment amidst deep practical, and moral, disagreement" (Adams 2024, 22) since, according to Adams's characterization of my view, "private moral judgment trumps the fact of disagreement and the need for politics" (Adams 2024, 23). The worry seems to be that, if whether we are obligated to obey the law 'because it is the law' depends on our private moral judgment and, realistically, we are bound to disagree on this matter, society will plunge into chaos. Different people will feel warranted in disobeying different laws (Adams 2024, 22).

This worry seems to me to be based on a misunderstanding. A non-realist does not think that whether we are obligated to obey the law depends on our private judgement. Instead, it depends on what morality requires. And if we are bound to disagree about the requirements of morality – without a publicly acceptable way of determining who is right – then we need to ask: how can we treat each other with respect, in a context where we disagree about what doing so, substantively, requires? Presumably, establishing some form of peaceful coexistence is a precondition for that, hence no non-realist in their right mind would suggest that citizens should just disobey the law whenever they think the law does not align with their preferred picture of morality. In short, Adams' worry about non-realists (like me) offering recipes for social collapse seems to me overblown.

One final point. Adams' own discussion suggests that there are two levels, when it comes to politics: a first-order level (level 1), in which people disagree about morality and politics, and a second-order level (level 2), in which we need to ask ourselves how to live together in light of the disagreement occurring at level 1. Work in political philosophy can also be pitched at two different levels. Some work will try to work out the 'correct view' at level 1, and some work will try to figure out the answer to the question occupying level 2. The contribution I offer in this book is pitched primarily at level 1. The realist channeled by Adams seems to proceed on the assumption that contributions at level 1 are pointless (or even dangerous), and that contributions at level 2 can only be offered from a realist perspective. Why this should be so, however, remains unclear to me.

References

Adams, N. P. 2024. "Moralism and Realism in Theorizing Social Norms." *Analyse & Kritik* 46 (1): 13—24.
Burman, Åsa. 2024. "Two Types of Social Norms." *Analyse & Kritik* 46 (1): 25—36.

Gilbert, Margaret. 2006. *A Theory of Political Obligation: Membership, Commitment, and the Bonds of Society*. New York: Oxford University Press.
Klosko, George. 2024. "Political Obligations and Respect for Social Norms." *Analyse & Kritik* 46 (1): 37—50.
Nieswandt, Katharina. 2018. "What is Conventionalism About Moral Rights and Duties?" *Australasian Journal of Philosophy* 97 (1): 15—28.
Nieswandt, Katharina. 2024. "Must I Honor Your Convictions? On Laura Valentini's Agency-Respect View." *Analyse & Kritik* 46 (1): 51—65.
Owens, David. 2006. "A Simple Theory of Promising." *Philosophical Review* 115 (1): 51—77.
Raz, Joseph. 1985. "Authority and Justification." *Philosophy & Public Affairs* 14 (1): 3—29.
Scanlon, Thomas M. 1990. "Promises and Practices." *Philosophy & Public Affairs* 19 (3): 199—226.
Stahl, Titus. 2024. "Social Norms and Obligation: Rescuing the Joint Commitment Account." *Analyse & Kritik* 46 (1): 67—83.
Valentini, Laura. 2023. *Morality and Socially Constructed Norms*. Oxford: Oxford University Press.
Valentini, Laura. 2024. "Do Socially Constructed Norms Have Moral Force? Précis to a Symposium." *Analyse & Kritik* 46 (1): 1—11.