



Titus Stahl\*

# Social Norms and Obligation: Rescuing the Joint Commitment Account

<https://doi.org/10.1515/auk-2024-2006>

**Abstract:** In *Morality and Socially Constructed Norms*, Laura Valentini argues that moral obligations to respect social norms can be explained without invoking the concept of ‘joint commitment.’ Her resulting account is, in one important sense, individualistic, and therefore struggles to account for widely held intuitions about the normative significance of social norms. I argue that we can rescue the notion of joint commitment from Valentini’s objections, and incorporate it into a version of her account that preserves its insights.

**Keywords:** joint commitment; obligation; social norm; individualism; social group

## 1 Introduction

In *Morality and Socially Constructed Norms*, Laura Valentini presents a groundbreaking and highly convincing argument for the idea that we have a pro tanto moral obligation to respect social norms that cannot be accounted for by the good consequences of complying with them or by their underlying moral quality.

In this article, I want to examine one particular aspect of Valentini’s argument more closely: Valentini argues that we can account for the existence of social norms in terms of individuals’ attitudes without introducing the concept of a ‘joint commitment’ and that the moral obligation to respect social norms is independent of whether those to whom it applies are members of the relevant community. I argue that Valentini’s account is in an important sense individualistic and therefore has trouble accounting for a range of plausible intuitions about the normative significance of social norms. I then develop a friendly proposal to add an unobjectionably minimal account of joint commitment to her argument that avoids these issues.

---

\*Corresponding author: Titus Stahl, Faculty of Philosophy, Groningen University, Groningen, The Netherlands, E-mail: [titus.stahl@rug.nl](mailto:titus.stahl@rug.nl)

Open Access. © 2024 the author(s), published by De Gruyter. This work is licensed under the Creative Commons Attribution 4.0 International License.

The paper proceeds as follows: In Section 2, I briefly describe why Valentini rejects the joint commitment account and what her alternative to it consists in. In Section 3, I show that this leads her into a form of individualism. In Section 4, I argue that this individualism makes a range of commonly shared intuitions seem mysterious. In Section 5, I present my friendly amendment and try to rescue the idea of joint commitment.

## 2 Why Not Joint Commitment?

The joint commitment view was most prominently formulated by Margaret Gilbert in response to what she views as the main challenges of individualist accounts of social norms (Gilbert 1989, 1999a, 2006). In short, Gilbert argues that we cannot explain why people's individual commitments to social norms generate an obligation, or a *standing* on the part of those enforcing the norm to exert "punitive pressure" on others, unless we assume that such standing and obligations arise through what she calls a *joint commitment*. Joint commitments are commitments that a group of people share not in an aggregative or summative way but as a "plural subject" (Gilbert 1989, 199) Plural subjects emerge through a process in which the involved individuals first bring their readiness to be jointly committed "out into the open" and then take each other up on that readiness, constituting a plural subject (Gilbert 1989, 215). Joint commitments are then to be understood as the commitments of such a plural subject. This explains, in particular, why individual members cannot simply rescind them unilaterally (Gilbert 1999b, 147), since we can only rescind our own commitments. Unlike personal commitments, though, joint commitments are not our own commitments in an individual sense, and we are therefore bound by them even though there are cases where we would rather not be.

It is this particular property of 'bindingness' on which Gilbert draws when she uses the joint commitment account to explain the possibility of social norms (Gilbert 1999a). Unlike conventionalist or 'shared intention' accounts, which only take individuals' intentions into account and which cannot explain why social norms are not fully up to us, understanding all social norms as rooted in joint commitments provides the necessary foundation for a form of social obligation that is not distinctively moral but shares with moral obligation an independence from the desires and attitudes of those to whom it applies, at least as long as we take them as individuals.

In her account of the normative force of social norms, Valentini rejects the joint commitment account both in terms of its descriptive adequacy and in terms of its capacity to account for our obligation to conform to social norms.

Valentini accepts that some social norms may be underpinned by joint commitments but argues that it is implausible to assume that all social norms are of this kind. Taking the social norm of queuing as an example, Valentini argues:

I find it odd [...] to suggest that we have jointly committed to upholding the queuing norm. [...] If I skipped the queue and someone ahead of me protested by saying “You are violating our joint commitment [...]”, I would respond that I do not know what commitment the person is referring to. (Valentini 2023, 77–78)

Valentini similarly argues that the joint commitment view fails to account for part of the normativity of social norms. First, we often have a moral obligation to respect social norms that are not plausibly rooted in any commitment we have joined, e.g. when visiting a foreign country (Valentini 2023, 111). Second, while the fact that we have intentionally joined a community that is committed to certain norms may sometimes give us additional reasons to comply with those norms, Valentini argues that it does not generate any additional moral obligation to do so (Valentini 2023, 112). Third, Valentini considers the argument that it is only through being part of a joint commitment that we are able to take up the ‘internal point of view,’ from which we can view these norms as binding on us as *ours*, rather than solely in virtue of their underlying moral rationale. However, Valentini argues that her preferred alternative – the agency respect view – allows us to understand how social norms can be morally binding in virtue of their social acceptance without requiring that this acceptance take the form of a joint commitment to which we are a party (Valentini 2023, 113).

As an alternative to the joint commitment account, Valentini argues for the *agential investment account* of social norms. According to this account, a social norm exists when it is widely and publicly accepted in a given context (Valentini 2023, 22). Accepting a rule involves two components: a belief that the rule ought to be followed, and a “robust intention that the rule function as a general standard of behaviour” (Valentini 2023, 26). This is compatible with the norm’s being generally disobeyed, since people do not always do what they believe they ought to or what they intend to do.

Valentini builds her explanation of the moral significance of such social rules – the agency respect view – on the agential investment account. According to the *agency respect view*, our obligation to respect others as persons entails a *pro tanto* obligation to respect the commitments that are central to their agency (provided they are morally permissible, authentic, and that respecting them is not too burdensome). Social norms, according to the agential investment account, are often an expression of people’s authentic commitments, and thus we must respect them (Valentini 2023, 97).

As Valentini explicitly holds, this obligation is independent of whether we are in any sense jointly committed to the norms involved (Valentini 2023, 111–13). Not only do social norms generally not involve joint commitments, but their moral significance is independent of whether those who stand under the obligation share any such commitment (jointly or otherwise).

Valentini holds that this account has the advantage of being explanatorily convincing since it does not introduce any new surprising normative principles beyond the widely accepted obligation to respect others (Valentini 2023, 101) and does not rely on any mysterious meta-normative claims. In addition, it fits the evidence for three of the central cases she discusses (Valentini 2023, 57–60): the *Traffic Light* case, which concerns the purported social norm in Germany of stopping when the traffic light is red even when no other vehicles are around and proceeding would not be dangerous; the *Barbecue* case, in which there is a social norm not to use other people's property without asking even if doing so would not harm them; and the *Non-Proceduralist President* case, in which a social norm explains why it is wrong to disregard legal procedures even for good purposes and as an exception. In all three cases, Valentini argues, we can understand the wrongness of breaching social norms independently of considering their moral quality or the effects of the breach, and she takes the agency respect view to offer the best explanation of why disobeying them is wrong.

### 3 Individualism and Parity between Personal Commitments and Social Norms

In this section, I argue that Valentini's account commits her to a specific form of individualism that directly entails that there is no significant normative difference between personal and socially shared commitments. This prepares the ground for my argument in the next section, where I show that this makes a number of common intuitions about socially shared norms seem unnecessarily mysterious.

The existence of social norms is, on Valentini's account, ultimately a matter of individuals having two kinds of attitudes, namely a belief (that a norm ought to be followed) and an intention (to see it followed). In terms of content, neither attitude necessarily refers to the intentions of others who follow the same norm; nor is either in any stronger sense shared (such that it is only held by people as members of a group). In other words, one could have the relevant combination of attitudes as an isolated individual without these attitudes thereby changing in terms of their character. What makes social norms social is (merely) that they are widely and publicly shared. Being widely shared is, I take it, just a matter of many people having the relevant set of attitudes, whereas their publicity is a matter of those

involved having another belief, namely a (true) belief to the effect that others have the same attitudes (including that very belief; see Valentini 2023, 29). It is, however, the individual belief and intention pair that accounts for the normative significance of social norms, since this pair is what we are obligated to respect.

Valentini's individualist account of social norms directly entails a form of *parity between personal commitments and social norms* concerning our obligations. There seems to be nothing particularly important, in principle, about social norms compared to individual commitments. Of course, our individual commitments are usually purely self-directed and do not refer to the actions of others, and thus they do not demand anything more of others, in terms of respect, than non-interference. Plausibly, however, people are not only committed to norms that they apply merely to themselves; they are sometimes individually committed, in the sense relevant to the agency respect view, to norms that they want to see followed more generally. Such personal commitments to norms – for example to norms regarding social behavior or the appreciation of cultural practices – can be central to who people take themselves to be. As long as such commitments are formed autonomously and authentically, and as long as they are not morally impermissible, it appears that Valentini's arguments seem to apply as much to such cases of personal commitment as they do to the case of social norms.

This leads to the following *parity claim*:

(*Parity between personal commitments and social norms*): There is no difference in normative significance between group  $G$ 's accepting a norm  $N$  as a social norm and each member of  $G$ 's being personally committed to  $N$ .

## 4 Six Problems for Individualism about Social Norms

The parity claim (which is directly entailed by the individualist analysis of social norms) faces difficulties when trying to account both for some of Valentini's own commitments and for common intuitions about social norms. In this section, I will discuss six problems for Valentini's view. As I will make clear, I do not assume that any of them are insurmountable or decisive, taken individually. However, I will try to show that, taken together, they suggest that there is something missing. In the next section, I will argue that there is an account of social norms that preserves the advantages of Valentini's account but is less individualistic and therefore has a better response to these problems.

- (1) The first problem is the *significance of publicity*. Valentini's account leaves it mysterious why the publicity condition for social norms matters, normatively

speaking. If what makes social norms morally significant is that they are a matter of people's robust individual commitments – which are perhaps particularly morally significant whenever they are held by *many* people – then it should not matter whether people know of each other that they have those commitments. Valentini argues that publicity explains the difference between “socially constructed” and “privately constructed” norms, such that we can say only of the former that “we share” them (Valentini 2023, 28–29). As I will argue later, however, there would seem to be nothing inconsistent about the idea of two people's personally being committed to a norm and knowing about each other's commitments without genuinely sharing a social norm.

More importantly, however, Valentini argues that publicity is a necessary condition for being accountable to one another, since we cannot hold each other accountable for failing to obey an esoteric principle (Valentini 2023, 29–30, 127). Accountability does not require publicity, however. For me to be able to follow, and to be accountable for following, a social norm constituted by the commitments of several people, it is only necessary that *I* know of *their* commitments (and in turn, they may need to be aware that I know this in order to hold me accountable); it is not necessary that *they* know about *each other's* commitments – or at least it is not clear why their knowledge about each other's commitments would make my obligation stronger than if that knowledge were only available to me.

On the basis of her account, Valentini cannot argue that publicity is a necessary condition for a norm's being a group norm in a stronger sense than merely being a matter of overlapping personal commitments. Intuitively, however, to say that ‘we share’ a norm is to say more than that.

- (2) A second problem concerns the idea that *communities can be the object of disrespect*. Valentini repeatedly suggests that, at least in some cases, by breaching a social norm we disrespect the community whose norm it is (see, e.g., Valentini 2023, 102) rather than any individual member. But if norms are merely a matter of the individual commitments of members, then it would be better to say that we violate an obligation to every member who has such a commitment. Similarly, Valentini argues that social norms enable communities to construct “a normative world” (Valentini 2023, 87) by being publicly committed to a norm. However, her analysis of norms leaves it mysterious why we ought to say that *communities* construct such worlds, rather than viewing them as the normative worlds of individuals that just happen to be identical. More generally, it also does not clarify whether the idea of a normative world carries any normative weight. Again, if we had the resources to say that norms can be norms of groups in a stronger sense than that allowed by Valentini's account, this would be less mysterious.

- (3) A third problem arises from the intuitively plausible idea that, from an outsider's perspective, there is a *disparity between overlapping personal commitments and genuine social norms*.

Imagine that I am invited to give a talk at another university. I learn that one of the colleagues at that university is strongly and robustly committed to the idea that speakers at their colloquium should wear ties. I, however, am strongly and robustly committed to cultivating an image of unconventional-ity and bohemian coolness that precludes wearing a tie. In this case, it seems (and I presume Valentini agrees) that I have no reason to take the colleague's personal commitment to be morally more important than my own.

Assume now that I learn of several colleagues at that university that they are each individually committed to this norm, unbeknownst to each other. Intuitively, it still seems that I have little reason to accommodate their commitments at the cost of my own. If I were indifferent about the matter, it would be wrong to frustrate their desires – but, as stipulated, I am not. In a third scenario, I come to know that they have learned of each other's personal commitments, such that they have now become public. They may be surprised and approve of others' happening to share their personal stance towards tie-wearing. Even then, however, it is unclear whether this generates a moral obligation on my part, since mere mutual knowledge of their relevant commitments does not seem to change the situation substantially.

Intuitively, all three scenarios differ quite substantially from the scenario that would arise if one of my hosts was to inform me that there is a *shared and jointly accepted* social norm at the university that prescribes wearing a tie. In that case, it seems as though there would at least be a weak obligation to conform to their social norm, even at the cost of disregarding my own personal commitments. This intuition again supports the idea that at least part of the normative force of social norms depends on their being shared in a stronger sense than being a matter of overlapping individual commitments.

- (4) A fourth and related problem concerns the *possibility of alienation* from a norm. By 'alienation' I mean a state of affairs where a person has a robust intention to enforce a norm, to follow it, and to accept criticism from others (on the condition that others have the same intention) but is not *personally* committed to the norm (in the sense of believing that it ought to be followed for reasons that are independent of the fact that it is currently accepted) and would not enforce or preserve the norm if others were to stop doing so. They might even prefer it if no one were committed to it (Brennan et al. 2013, 33–34). I assume that many of us are alienated, in this sense, from at least some of our social norms and that there are perhaps cases where everyone involved is alienated. At my university, for example, there is a norm to the effect that

during a PhD defense, candidates must address each committee member as ‘highly learned opponent.’ Personally, I find this norm slightly silly; if no one thought that it ought to be followed, I certainly would not do anything about it. I suspect that many, perhaps all, of my colleagues feel the same. At the same time, it seems clear to me that this is still a genuinely shared norm in the sense that most faculty members believe that candidates should follow it as long as it remains in force and continue to instruct them to do so.

Valentini’s account also supports this conclusion, it seems to me. Even if people are alienated from a norm, it can still be said to be a genuine social norm, as long as they believe that it ought to be followed (conditional on other people’s believing this as well) and as long as they intend to see it being followed (conditional on other people’s having that intention). Such a norm also seems to generate obligations for those subject to it, since choosing not to follow it seems to be (at least minimally) disrespectful of the fact that there is a community that has adopted this as their norm and has assigned to some of its members a right to be addressed in this way.

On Valentini’s account, however, it is unclear why there is any moral obligation to obey it, since the conditional commitments that people have are not evidence that the norm being followed is important to their individual agency – if anything, their alienation provides evidence to the contrary. Thus, there would seem to be cases where an obligation to obey a social norm exists even though not abiding by that norm would not be disrespectful to any individual’s personal commitments.

- (5) A fifth problem concerns the question of the significance of ‘*our own norms*.’ This is the question whether there is anything particularly wrong with failing to obey social norms that one endorses oneself. This seems to be ruled out by the parity claim since, if disrespecting the personal commitments one shares with others is just as wrong as disrespecting those one does not share, the same should be true of social norms.

Valentini considers this question and argues that, for the most part, there is no real difference between moral obligations for “insiders” and “outsiders” (Valentini 2023, 111–13). She grants that “one may have additional reasons, beyond agency respect, to act as the norms of one’s own community prescribe” (Valentini 2023, 112) but argues that “it is not clear that [the relevant commitments] give rise to any additional moral obligation to obey” (ibid.). She instead argues that affirming norms may be valuable for agents, but not morally obligatory, in “the way in which appreciation of the arts is valuable but not obligatory” (ibid.).

This seems to entail that when we fail to obey norms that we endorse as part of a community, we may fail to properly react to some reasons. But we are not

blameworthy in any special way, and no one has any particular standing to demand compliance beyond what would be the case for norms to which we relate as outsiders (just as we are not blameworthy for not appreciating the arts and no one has standing to demand that we do so).

Intuitively, however, it seems that when we violate norms that we genuinely share with others, this entitles those others to demand our compliance in a particular way (for this feature, see Gilbert's discussion of Hart in Gilbert 1999a, 150–51). Imagine that the members of a family are committed to the norm of calling each other on their birthdays. They may never have explicitly agreed to follow this rule, but it is understood between them that all of them believe they ought to follow it, that they robustly intend to see it followed, that they use it as a standard of evaluation, and so forth. When a family member fails to obey the norm, it would seem that it is not only the person whose birthday was forgotten but all other members of the family who have special standing to criticize that failure. They might say, for example: "You violated a norm that we all accept" or "our norm" (Gilbert 2001, 43). There is nothing strange about saying, in such a case, that the norm violator has not merely wronged the person who was entitled to receive a call on their birthday but violated a commitment that they hold *as a member* of the group.

The same appears to be true in the case of larger groups, where there is a clear mutual commitment to follow the relevant norm. Assume, for example, that in a specific academic discipline there is a norm of thanking every single person with whom one has previously discussed a subject when one publishes an article on that subject. Again, if such a norm is jointly accepted in a community in the sense that people believe it ought to be followed, intend to follow and enforce it, and use and accept it as a standard of criticism, then it seems that they may criticize norm violators for having failed to abide by the shared norm of the community.

Does this standing to criticize derive from the fact that the norm violator has acted against the strong, identity-defining commitments of the person who criticizes them? It seems to me that this is not sufficient to account for the full force of such criticism.

Imagine that I am committed, in the relevant strong sense, to the norm that people should always call their family members on their birthdays. This, in itself, does not seem to give me any particular standing to criticize others for failing to call their family members, even if the same norm is accepted in their families. You may accept that I am strongly committed to seeing that norm obeyed throughout society, and you may even accept that you have a weak, *pro tanto* obligation not to unnecessarily frustrate my intention to see that norm followed. However, my entitlement to criticize is certainly much

weaker than the entitlement of *your* family members. In particular, it would be strange if I were to say: “You violated *our* norm.” There is, strictly speaking, no real *we* here (see Gilbert 1989, 167–203).

Similarly, in the academic case, if people outside the relevant academic community were for some reason to form a commitment to the idea that academics should follow the norm of always thanking all interlocutors, we would perhaps have a weak obligation to humor them so as not to undermine their commitments unnecessarily. But clearly, if outsiders were to claim to have standing to criticize us for failing to follow *our* shared norm, then ‘Who are you, again?’ would be an entirely appropriate reaction.

We can therefore conclude that the special standing of community members is not explained entirely by the fact that they are personally committed to the norm, since in the cases described outsiders can be so committed. It is also not explained by the fact that they are designated by the norm as being entitled to a certain kind of treatment. The special standing of insiders seems to persist in cases where they are not beneficiaries of the prescribed action. Similarly, it is not necessary for those who criticize us to see themselves as potentially entitled to similar treatment by others. We can easily modify the cases such that there is a norm of calling children on their birthdays and one is criticized by a co-parent, or where only PhD students need to thank others but are criticized by senior faculty members. Even in these cases, the intuition persists that genuinely sharing a norm with others gives those others special standing. As such standing only exists in the case of shared but not overlapping personal commitments, there must be a difference between merely aggregatively (even if publicly) shared personal commitments that do not create a ‘we’ and genuinely shared social norms in communities.

- (6) This leads to a sixth problem that is closely related and concerns the idea of *separate communities with the same norms* (Noyes et al. 2023, 153; Stahl 2022). If we assume that a social norm’s existence is a matter of individual belief and individual intention, it would appear to be impossible for two groups of people to happen to accept the same norm without genuinely *sharing* that norm – if every member of the two groups has the relevant belief and intention, and if this is publicly known, then all of Valentini’s conditions for the existence of a social norm are met.

It is entirely plausible, however, to suppose that different groups might separately accept the same norms. Not only might members of separate groups accept the same norms, in each case directed at the other members of their own group, but two groups might separately endorse norms that prescribe behavior that they also expect members of the other group to obey. Assume, for instance, that both a religious group and a secular temperance movement

endorse the norm that no one should consume alcohol (including members of the other group). It would be counterintuitive to say that there is a genuinely *shared* norm in play here. This is because members of each group would view members of the other group as outsiders who have no special standing to demand norm compliance from them.

Of course, one could argue that this is a matter of two groups endorsing the same standard for different reasons and that a necessary condition for a group's sharing a norm is that all members must form their belief and intention on the basis of the same underlying reasons. However, not only would this be a departure from Valentini's account, which contains no such requirement, but it would also be implausible, since many groups seem to have genuinely shared norms even though their members do not endorse them for the same reasons.

As indicated above, these six problems – *the significance of publicity, community disrespect, disparity, alienation, insider standing, and separate communities* – do not amount to a conclusive objection to Valentini's account of social norms. Taken together, however, they indicate that endorsing this account comes with certain costs. These problems are sufficient to at least justify considering whether there are alternative accounts of social norms that avoid these problems and preserve the insights that Valentini offers.

## 5 Rescuing Joint Commitment

In this section, I will develop a version of the joint commitment view that can be seen as a non-individualist extension of Valentini's agency investment account. I will argue that this revised account avoids some of the features of the joint commitment view that she describes, with good reason, as objectionable and that it is better able to account for some of the intuitions I have just described.

To recap the discussion in Section 1, Valentini's arguments against the joint commitment view are as follows: (1) We have an obligation to obey social norms in cases where we clearly are not party to any joint commitment, e.g. when we visit foreign countries. (2) It is often implausible to refer to joint commitments when discussing social norms. (3) Joint commitments do not give rise to any independent moral obligations.

In what follows, I do not want to deny (1). Valentini convincingly argues that we have a pro tanto moral obligation to respect the agency-defining commitments of others, and that social norms that we ourselves do not accept are often backed by such commitments. From this, she concludes that an explanation of the moral force of social norms need not refer to joint commitments. However, as I argued in Section 3, not only do we intuitively consider there to be a difference between

our obligation to respect the personal commitments of others and the obligation to respect the joint commitments of others, but there is an intuitive difference between a relatively weak obligation to respect social norms to which we relate as outsiders and the stronger obligation that applies to insiders. Even if Valentini is correct that referring to joint commitments is not necessary to explain certain aspects of the obligations we have regarding social norms accepted by others, this does not entail that there are no other forms of social or moral obligation that could be explained by a joint commitment view. And intuitively, at least, there seem to be cases where joint commitments matter.

Regarding (2) – the claim that it is often implausible to refer to joint commitments when discussing social norms – Valentini argues that when we ask ourselves why we should conform to the norm of queuing, for example, it is unnatural to assume that there is any joint commitment.

Of course, very few people would use the terminology of ‘joint commitments’ when describing themselves and their obligations. However, it does not seem unnatural to draw a distinction between cases where we are only confronted with the commitments and expectations of others, which we do not share, and commitments that are in some stronger sense *ours*. In the example of queuing, it does not feel unnatural to say: ‘You are violating a norm that we follow here,’ or ‘This is our norm here.’ To say this is to say more than that there are many individuals around who are strongly committed to everyone’s following the norm; it is to say that this is a norm that is collectively (dare one say jointly?) accepted by the community that makes up the context of the interaction.

Part of what makes it implausible to view this intuitive idea as supporting Gilbert’s joint commitment account is that Gilbert has a rather narrow view of the conditions for the existence of joint commitments, which can only come into being, on her view, if all members of the relevant group each announce their ‘readiness’ to be jointly committed (Gilbert 1989, 185–99). Not only does this leave it mysterious how this readiness is transformed into an actual joint commitment (Stahl 2013), but it also gives the impression that joint commitments can only emerge through some kind of quasi-contractual procedure in which all consent in some way or another.

In her application of this idea to larger groups, such as political communities, Gilbert argues that such awareness of mutual readiness can also slowly emerge in a larger society through a general awareness of the fact that unspecified others are ready to commit (Gilbert 1989, 215; 2006, 199). Nevertheless, there still seem to be two steps involved in this picture: becoming aware that everyone is ready to be jointly committed and then actually being so committed. This is, of course, not very plausible when we think about cases such as queuing, where we cannot be said to have been ‘ready to commit’ and then, in another step, to have ‘entered’ into the relevant joint commitment.

This may well explain Valentini's reluctance to think of social norms as a matter of joint commitment. However, this particular feature appears to be more of an artifact of Gilbert's account than a necessary component of the idea of joint commitment itself.

Consider the following alternative proposal:

JOINT COMMITMENT: A group  $G$  is jointly committed to a norm  $N$  if and only if

- (a) every member of  $G$  believes that  $N$  ought to be followed (conditional, perhaps, on everyone else's believing it);
- (b) every member of  $G$  robustly intends that  $N$  will function as a standard;
- (c) every member of  $G$  understands themselves to be required to accept criticism regarding their compliance with (or enforcement of)  $N$  by other members as *representatives of  $G$* , and understands themselves as entitled to criticize other members *in one's role as a representative of  $G$*  if they fail to comply with (or enforce)  $N$ . (Stahl 2022, 207–8).

Even though I am not certain that we should see (a) and (b) as necessary conditions for the existence of a social norm, I want to stay as close as possible to Valentini's account to preserve its attractive features. Therefore, this proposal adds another necessary condition that introduces other-directed attitudes, that is, practical attitudes by which one accepts the authority of *specific others* regarding one's compliance with  $N$  as long as these others evaluate one as representatives of  $G$  (and takes oneself to be entitled to that acceptance). *To act as a representative of a group* means to evaluate others not on the basis of one's personal preferences or understandings, but on the basis of an interpretation of what norm the group has accepted that is responsive to challenges from others (Stahl 2022, 150). While this account remains close to Hart's (1994; see also Brennan et al. 2013) account insofar it explains the existence of norms in terms of normative attitudes, these additional features allow it to escape Gilbert's criticism of Hart's individualism (Gilbert 1999a).

I think it is plausible to say that any group that meets these conditions can be said to display genuinely *joint* commitments, since no person in  $G$  is jointly committed to  $N$ , unless all others are also so committed. Joint commitment is thus more than just a matter of coinciding personal commitments.

This proposal clearly avoids the idea that in order to be jointly committed to a norm, one must first be ready to join a commitment or advertise one's readiness. People often acquire the necessary beliefs, intentions, and understandings as a matter of being socialized in a community. Therefore, there appears to be nothing odd about saying of a group of strangers at a bus stop that they are jointly committed, if this merely means that they have these attitudes towards one another.

At the same time, the proposal also lays the groundwork for rejecting the idea that there is parity between merely personal commitments (in which members of a group meet conditions (a) and (b) but not (c)) and genuinely joint commitments.

This difference also sheds helpful light on the issue of *separate groups*. Two groups can be committed to the same standard of behavior (even a standard they want to see followed by the other group) without being jointly committed, as long as the members of each group only accept the authority of members of their *own* group regarding evaluations of their behavior.

It also helps us to understand the issue of the special standing of fellow group members and thus to make sense of the particular forms of criticism that are only appropriate in the case of joint commitments. In the case of queuing at the bus stop, assume that we are party to a joint commitment in this sense. This means that we accept that others are entitled to criticize us should we fail to comply. In this case, when someone says ‘This is what we do here,’ the *we* successfully includes us as members of a group of people all of whom stand in this relation to each other. It expresses the speaker’s (correct) assumption that we accept them as entitled to voice such criticism. If we spell out the idea of ‘joint commitment’ in this way, it would seem natural to say that it is indeed in play in everyday cases. Moreover, this approach is well suited to making sense of our intuition that, when it comes to others’ standing to criticize us, whether we are party to such a joint commitment is indeed relevant.

The account also helps us to make sense of the possibility of alienation. In cases where we are alienated from a norm, we continue to accept others as entitled to demand that we obey it or enforce it, and we may even be inclined to demand such actions of others ourselves. Nevertheless, we only believe that it ought to be obeyed in the minimal sense that we think that, *given these mutual attitudes*, people should obey it. We are not independently personally committed to it.

The significance of publicity is equally easily explained: It is difficult to see how we could consider each other as entitled and obligated in this way in the absence of some understanding that others have the same attitudes. Therefore, cases in which people are generally unaware of each other’s commitments cannot ever be cases of genuinely shared norms. It also explains Valentini’s intuition that publicity is a condition for mutual accountability – in order for a norm to be genuinely shared, there first needs to be general acceptance of mutual accountability, and this acceptance then allows for an internal group perspective from which stronger commitments become intelligible.

This leaves open two questions: first, whether the joint commitment account of social norms also accounts for stronger obligations to abide by the shared norms of other communities, beyond the agency respect principle, and second, whether Valentini is right to say that joint commitments do not create any moral obligations.

Both questions can be answered together. The key to answering them is the idea that joint commitments create social obligations by virtue of the fact that having a joint commitment implies an understanding of oneself as a member of a community, such that membership implies an obligation towards that community. As Southwood (2011, 790) discusses, one may object that this allows for a form of ‘objectionable bootstrapping’ as it seems as though the obligation comes into existence merely in virtue of people’s accepting it. Many practices come with a specific understanding of the reasons for accepting such obligations – reasons that only emerge if one accepts norms regarding the proper appreciation of values that are internal to the practice in question. While we may be unable to account for the grounds of such obligations from the outside, accepting them may be reasonable, even inevitable, from the inside perspective that is itself mandated by the rules of such a practice. In Valentini’s own terms, practices bring ‘normative worlds’ into being – but only when they are governed by joint commitments.

Such obligations that one can only appreciate on the basis of reasons that are internal to a practice are not moral obligations. This is because obligations can only be characterized as moral, i.e. categorically binding, if anyone can understand them as binding, independently of the internal perspective of any particular social practice of which they may or may not be a member (Brennan et al. 2013, chap. 4). However, practice-internal obligations are clearly *morally significant*, from both the outsider and the insider perspective. This is because we can extend the agency respect principle to joint commitments when we appreciate that people’s agency is often not only shaped by personal commitments but sustained by and dependent on participation in joint commitments and the practices they enable (Anderson 2010, 7–12). If it is central to my agency to be committed to the identity of an academic, and I cannot have such an identity apart from there being joint commitments that bring an academic community into being, then the moral obligation to respect my agency also includes the obligation to respect the joint commitments that come with my membership in this community and that enable me to maintain a specific agential identity.

The obligation to respect joint commitments is typically morally weightier than the obligation to respect another person’s personal commitments, for two reasons. First, when we disrespect a community’s joint commitments, we typically disrespect the conditions of agency not only of those with whom we are interacting but also of all other members of the group whose commitments they are. Second, and more importantly, disrespecting the personal commitments of another person may make it psychologically harder for her to sustain them. Clearly, however, it does not affect the availability of the option of her being so committed. By contrast, when community-constituting joint commitments are undermined and made harder for people to sustain, communities can break down. Such a breakdown

often removes the option of participating in the joint commitment. Once a joint commitment has been undermined and has disappeared, it is often difficult or even impossible to bring it into being again. Therefore, the moral harm that results from disrespecting joint commitments is often much greater than the moral harm caused by disrespecting personal projects.

Of course, not all joint commitments are authentic, important and worthy of respect. In addition, the obligation to respect joint commitments can be outweighed by the fact that compliance is too burdensome or by a justified allegiance to an incompatible, more weighty personal project. As a general principle, however, we can safely assume that the obligation to respect community-constituting and identity-defining joint commitments is typically weightier than the obligation to respect personal commitments, which explains the typical disparity, for outsiders, between obligations regarding social norms and obligations regarding personal commitments.

For insiders, the same considerations explain why we often have an even stronger *pro tanto* obligation not to violate the social norms of *our* communities. It is not merely a commendable attitude to do so (as Valentini argues). The fact that others rely on *us, as co-members*, to sustain the social norms of a practice that gives them access to internal reasons the appreciation of which is central to their identity and agency makes it *pro tanto* morally wrong to deny them our cooperation in this project.

This does not mean, of course, that one may never abandon a joint commitment when one has good reason to do so; social norms can become problematic for a number of reasons, and continued participation in such a commitment can undermine one's pursuit of other projects that have become central to who one is. While outsiders merely have a *pro tanto* obligation to respect our joint commitments, however, insiders have a (defeasible) *pro tanto* obligation to also *sustain* them as long as other members depend on their cooperation for the continued existence of practices that are essential to who they are. For this reason, the proposed joint commitment view can explain both the disparity between personal commitments and social norms from the perspective of outsiders and the special standing of insiders in terms of their moral significance.

## 6 Conclusion

I have argued that while Valentini's agency respect theory is plausible and convincing, its underlying individualist account of social norms makes some of our intuitions unnecessarily mysterious. Amending her theory of social norms with an unobjectionably minimal concept of joint commitment avoids this individualism

and preserves the attractive features of her account. Since the principle that explains the moral significance of social norms – the agency respect principle – remains unchanged in this revised account, the grounds of our obligations remain as unsurprising and non-mysterious as before. It also continues to be supported by all the evidence that Valentini marshals. Clearly, the moral obligation to respect German traffic rules, norms about how to treat other people's property and the procedures of our political institutions is, if anything, even easier to understand if we assume that the relevant norms are the objects of joint commitments. Accepting Valentini's account therefore does not force us to abandon the idea that social norms are a matter of joint commitments. On the contrary, her account offers rich support for this idea.

## References

- Anderson, Elizabeth. 2010. *Value in Ethics and Economics*. Cambridge: Harvard Univ. Press.
- Brennan, Geoffrey, Lina Eriksson, Robert E. Goodin, and Nicholas Southwood. 2013. *Explaining Norms*. Oxford: Oxford University Press.
- Gilbert, Margaret. 1989. *On Social Facts*. London: Routledge.
- Gilbert, Margaret. 1999a. "Social Rules: Some Problems for Hart's Account, and an Alternative Proposal." *Law and Philosophy* 18: 141–71.
- Gilbert, Margaret. 1999b. "Obligation and Joint Commitment." *Utilitas* 11 (2): 143–63.
- Gilbert, Margaret. 2001. "Social Rules as Plural Subject Phenomena." In *On the Nature of Social and Institutional Reality*, edited by Eerik Lagerspetz, Heikki Ikäheimo, and Jussi Kotkavirta, 39. Jyväskylä: SoPhi/University of Jyväskylä Press.
- Gilbert, Margaret. 2006. *A Theory of Political Obligation: Membership, Commitment, and the Bonds of Society*. Oxford, New York: Oxford University Press.
- Hart, H. L. A. 1994. *The Concept of Law*, 2nd ed. Oxford: Oxford University Press.
- Noyes, Alexander, Frank C. Keil, Yarrow Dunham, and Katherine Ritchie. 2023. "Same People, Different Group: Social Structures Are a Central Component of Group Concepts." *Cognition* 240: 105567.
- Southwood, Nicholas. 2011. "The Moral/Conventional Distinction." *Mind* 120 (479): 761–802.
- Stahl, Titus. 2013. "The Conditions of Collectivity. Joint Commitments and the Shared Norms of Membership." In *Institutions, Emotions, and Group Agents*, edited by Hans-Bernhard Schmid, and Anita Konzelmann Ziv, 229–44. Dordrecht: Springer.
- Stahl, Titus. 2022. *Immanent Critique. Essex Studies in Critical Theory*. Lanham: Rowman & Littlefield International.
- Valentini, Laura. 2023. *Morality and Socially Constructed Norms*. Oxford: Oxford University Press.