
JACOB HEIM
University of California, Irvine
jheim@uci.edu

COMMITMENTS IN GROUPS AND COMMITMENTS OF GROUPS

abstract

I argue that a group can have normative commitments, and that the commitment of a group is not merely a sum or aggregate of the commitments of individual group members. I begin with a set of simple cases which illustrate two structurally different ways that group commitments can go wrong. These two kinds of potential failure correspond to two different levels of commitment: one at the individual level, owed to the other group members, and one at the group level, which the group as a single body owes either to itself or to some third party. I distinguish the content of a commitment (what must be done for the commitment to be fulfilled) from the holder of that commitment: the party to whom the content is owed. I then discuss examples which support the two-level view of group commitment and show that, even when individual-level and group-level commitments have the same content, they are understood to have different holders. Finally I return to my original cases and argue that a two-level structure of group commitment allows us to make sense of the problems that occur in them.

keywords

social ontology, joint commitment, Edward Snowden

When we talk about the commitments that exist in a group setting, it seems natural to ask either one or both of the following questions: (1) What is the group committed to; and (2) What are the group's members committed to? I will be arguing that a complete answer to (1) will involve more than a mere list or aggregation of the answers to (2). To put it another way, the commitments of a group cannot be fully reduced to the commitments that exist *within* that group. When we closely examine the ways that group membership constrains our activities, we find that these constraints entail the existence of group-level commitments that can only be honored by keeping them distinct from our commitments as individuals.

While I will not be assuming some particular theory of group commitments from the outset, the considerations that I raise will pose more of a challenge for some views than for others. As should be clear by the end of the paper, accounting for two distinct levels of joint commitment will be much easier within a framework like Margaret Gilbert's "plural subject" theory than it would be for a more reductive, individualistic view.¹ If we want to take seriously the standards to which we hold one another in practice, we are going to need to make room in our social ontology for the existence of group-level commitments that cannot be reduced to anything on the individual level.

1. Group and personal failures

Though we talk easily enough about what this or that group is committed to, we might think that such talk is merely a convenient shorthand for what the *members* of the group are committed to.² After all, groups are made up of individuals, and what a group does is in some sense done by those individuals: when someone says that the team wears a yellow uniform, we might interpret her as saying that the *members* of the team wear yellow uniforms. If properties of groups can be fully analysed into properties of individuals, then we should expect that any case in which a group fails to fulfil its commitment can be fully described in terms of individual commitments and individual failures to fulfil them. We will see whether this turns out to be the case.

Let us begin by looking at a case of individual commitment. To avoid unnecessary complications, we will consider the simplest kind of situation in which a commitment can arise: one in which there is only a single agent, making a commitment to himself.

1 For the initial statement of her view, see Gilbert (1989). As far as I can work out, everything I say here is consistent with Gilbert's overall theory, but does not presuppose it.

2 Bratman (1999), Brewer (2003), and Wilkins (2002) all seem to be thinking along these lines.

Case 1: I am sitting on my couch watching TV. It occurs to me that I need to go to the store before it closes, and I decide to go as soon as this show is over. When the show ends, I get up and go to the store.

One might wonder whether a case like this counts as a commitment at all. I would answer that there are (at least) two senses of the term *commitment*, which refer to related but distinct concepts. One, the stronger sense, is what we mean when we say, for example, that Martin Luther King, Jr. was *committed* to the cause of racial equality.³ The other, weaker sense is the one that I am using here. To be committed in this sense does not require any connection to one's deep-seated beliefs or goals; if I tell you that I will meet you at 2:00 then I have a commitment to do so, regardless of how much or little I care about it. This is the sense of commitment that we generally use when we talk about what so-and-so is committed to *doing*. Importantly, one can have a commitment in this sense without having made that commitment to anyone else: if I decide to do something then I am committed to do it.⁴ My deciding to do *x* is in this way normatively equivalent to telling you that I will do *x*, except that in the former case I owe the performance of *x* to myself rather than to you. In the terminology I will be introducing later, I am the *holder* of the former commitment, and you are the holder of the latter.

Many of the commitments discussed throughout this paper could also be properly described as obligations. I will stick to the term *commitment*, however, as it seems to be the more general concept, and it is at that level of generality that my arguments (and conclusions) are intended to operate. While I see no reason not to assume that all obligations are commitments of some kind,⁵ there do seem to be commitments (of a kind relevant to the topic of this paper) that are not, at least obviously, also obligations. In particular, it would be slightly odd to say that in Case 1 I have an *obligation* to myself.

Case 2: I am sitting on my couch watching TV, and I decide to go to the store after this show. But when the show ends, I do not move. I continue to watch TV until long after the store has closed.

Here I have failed in my commitment to myself. I would be entitled to rebuke myself in much the same way that you could rebuke me if I had told you I would go to the store for some purpose of yours.

Now let us move from a simple individual case to a simple group case.

Case 3: You, me, Amy, and Chad are sitting on the couch watching TV. It occurs to one of us that we need to go to the store before it closes, and after a brief discussion we agree to go together as soon as this show ends. But after the show, instead of moving, we continue to sit there.

Something has clearly gone wrong here. Perhaps each of us is waiting for the other to make a move, or perhaps we have all gotten distracted. But for whatever reason, we have made an agreement – the interpersonal analogue of a personal decision – and then failed to comply with it. The failure here is comparable to the one in Case 2. But not all group failures align so neatly with the individual case. Consider the following:

3 This example is borrowed from Talbot Brewer (Brewer 2003, p. 562).

4 See Gilbert 2006 (pp. 127-134) for her argument that a decision is a form of commitment to one's self.

5 Since nothing in my argument hangs on this one way or the other, I will not bother arguing for it.

Case 4: You, me, Amy, and Chad are watching TV. We agree to go to the store together as soon as this show is over. When the show ends, you, Amy, and Chad get up and move towards the door while I continue to sit there. You remind me of our agreement, and I vaguely reassure you that I intend to get up any minute now. After a while the three of you grow tired of waiting, and go to the store without me.

Once again, there has obviously been a misfire of some kind. But this is the kind of misfire that can only happen in a group setting; a single person cannot simultaneously perform and fail to perform the same action.

Pre-theoretically, it seems that the most natural way to describe what has gone wrong differs significantly from Case 3 to Case 4. In Case 3, it may be appropriate in some sense to say that each of us failed the others: after all, none of us did what we told the others we would do. But that does not seem to get at the heart of the problem. No one of us was under any particular obligation to get up first, so none of us is in a position to blame our own inaction on the inaction of another. Each of us has a decent excuse, even if it is not quite a good one, for not getting up first – after all, why should it have to be *me*? What complicates matters is that the rest of us have an identical excuse. One could imagine the four of us discussing this very issue, after the fact, and finally coming to the conclusion that the most appropriate thing we could say about it is simply that *we* failed. There need not be anything more to say as far as who failed whom on an individual level. Indeed, we may come to realise that to say anything at all, on that level, would be unfair.

Case 4 is quite different. Strictly speaking, it would still be correct to say that we did not do what we said we would; the plan was for the four of us to go to the store together, and that did not happen. The responsibility for that, which could not be assigned in Case 3, now falls squarely on me. But not only is it my fault that *we* fell short of *our* plan. Additionally, and more importantly, *I* failed the three of you by not doing what *I* agreed to.

What we have, it seems, are two different and apparently incommensurable ways in which group commitments can fail. In one, a group fails as a group, whether or not responsibility for that failure can be traced back to particular members. In the other, an individual group member fails as a group member, whether or not the group itself ends up failing as a result. Both kinds of failure occur regularly in our social lives. Of course, they can and do coincide often enough. But there are many cases in which they do not, and in practice we readily distinguish between them in the ways that we respond. Thus a satisfying account of group commitments should be able to explain failures of both kinds in a way that does justice to our ordinary way of thinking about them.

- 2. A distinction about commitment**
- Given that our task is to explain two distinct kinds of commitment failure, a natural first step would be to try to identify two distinct ways of being committed, to which the failures might correspond. One of the failures, Case 4, is fairly straightforward. It involves a commitment that arises in the context of a group, but as we have seen, the nature of the failure suggests that the relevant commitment belongs to a particular individual only. If failures like Case 4 were all we needed to account for, we might be able to get by with personal commitments alone. The other failure calls for a more nuanced approach, however, as Case 3 certainly *appears* to be a failure by the group as such, which in turn suggests a commitment by the group as such. This appearance must either be explained, or explained away. In the absence of compelling reasons to view group commitments as personal commitments in disguise, we ought to at least take seriously the possibility that a group can operate as a unified normative entity, capable of having its own rights and obligations. Once we have admitted the possibility that our everyday talk of group commitments might

refer to something real and irreducible, we can begin to ask meaningful questions about the ways in which a particular group member can relate to a commitment of the group. In order to pursue these questions, I will set aside the question of how joint commitments are formed and will focus instead on some of the ways that such commitments work themselves out in situations where they are understood to exist.

When a group of people undertake a joint commitment, the group acquires a commitment which is different in kind from the aggregated commitments of the individual group members. This difference in *kind* of commitment is one that holds even when the *content* of the group-level commitment does not contain anything over and above the sum of personal commitments within the group. By “content” I mean the actions to be performed, or the conditions that must obtain, in order for the commitment to be considered fully met. Distinguish this from what I will call, for lack of a better term, the *holder* of the commitment: the party to whom the content of the commitment is owed, and (in general) whose right it is to declare the commitment satisfied or not.

Commitments can be differentiated with respect to either one of these features, even if they are identical with respect to the other. If I promise you that I will wash the car and walk the dog, I have two distinct commitments, even though they have the same holder, because they have different content. If I tell you that I will make us something good to eat, and I tell myself that I will make us something good to eat, there are two commitments as well, because they have different holders despite identical content. Admittedly, this is less obvious than in the case of differing content, but consider: if I make food that you greatly enjoy but that I find revolting, you will be satisfied that I have done what I said I would do, while I would not think the same. In this case I have met my commitment to you, but not to myself.⁶

It may be worth mentioning at this point that there is a common type of occurrence in which a group can have a commitment that differs from the commitments of its members, not just in terms of holder, but in terms of content as well. This is the case whenever a group has taken on a commitment but has not yet determined which group members will be responsible for which components of the larger commitment. Opponents of plural subject theory might want to explain this phenomenon away by arguing that each group member does have a personal commitment to do her part in the larger project, even before it has been determined what that part is, and that necessarily the sum of the as-yet-undefined personal commitments will add up to the group’s commitment. I do not find this explanation satisfying, partly because, as I understand it, this is just what it means to be party to a commitment larger than one’s own. Still, our imagined opponent is onto something, in that whatever actions the group is committed to do must ultimately be made up of the actions of its members. An orchestra plays a symphony, but each note of that symphony is played by a particular member of the orchestra. But even when the content of the group’s commitment is nothing more than the aggregated content of the individual commitments – indeed, even if we were to grant that such is always the case – a group-level commitment nevertheless differs from the related individual commitments in that it has a different holder. An example will help to illustrate.

If a group of friends go out to a pizzeria and order a large pizza together, the cost of the pizza creates parallel commitments on two levels. As a body, the group sitting at that table has a commitment to pay the full cost of the pizza, while each person at the table has a commitment to pay a fair portion of that bill. However the individual shares get divided, they will add

**3.
Two levels of
commitment in
practice**

⁶ In Roth’s terminology, I would have fulfilled my *contralateral* commitment but not my *ipsilateral* commitment (Roth 2004).

up to the cost of the pizza which the group is committed to pay for. So the content of the commitments at the two levels are identical.

They do not have the same holder, though. The individual friends each owe it *to each other* to pay a fair share, and it is only the group as a whole that owes it to the restaurant to pay for the whole pizza. If one person is bent on paying less than her share, it is the other friends at the table who have standing to rebuke her; the waiter does not. As long as the full amount is paid, the waiter has no standing to complain about who pays what. And if the stingy friend wants to defend her position – perhaps she ate less pizza than the others and so feels entitled to pay less – she would naturally address herself to the other members of her party rather than to the waiter.

Suppose that she does otherwise: she ignores her friends and explains herself to the waiter directly, insisting that she has contributed a reasonable share and that someone else is going to have to make up the difference. This would not be merely odd, it would be out of line. It would be natural, in such an instance, for one of her friends to protest, “Don’t bring *him* into this!”. Even the waiter himself could say as much, politely insisting that they settle it among themselves. We would normally regard this as the proper way for him to respond: it is not his business whether a particular group member pays a fair portion of a shared bill. That is a commitment internal to the group, not a commitment between individual group members and the restaurant. She cannot individually make good her commitment with the restaurant, because she has no such commitment. Among the set of her commitments, there is not one whose content is that she pays for her share of the pizza and whose holder is the restaurant. In the extreme case, if she were to simply hand five dollars to the waiter and firmly insist that she will not pay another penny, this would rightly be regarded as a serious offense – not against the restaurant, but against the other members of her party and against the group that they constitute. By trying to isolate a single share of what is properly a group obligation, and presuming to settle it individually with the holder of that group obligation, she has directly violated one of the most basic norms of group membership.

Similar cases are easy to come by. You hire a construction crew to build a shed, and the finished product falls over the first time it rains. Naturally, you demand a refund. It would be outrageous for one of the workers to come to you and say that *his* work on the shed was not the cause of the collapse, and so you are not entitled to a refund on *his* share of the money. He cannot claim to have fulfilled his individual commitment to you, because he has no such commitment. As far as you are concerned, there was only ever a *group* commitment to build a complete shed, and in that commitment the entire group failed as one body. Whose handiwork is to blame, and who ought to be penalized and how, is a conversation for *them* to have.

What these examples show is that we do in fact distinguish, in ordinary practice, between individual-level and group-level commitments. There are real social consequences for failing to respect these distinctions. Commitments on the individual level are held by the group, and are owed to one’s fellow group members. Commitments at the group level are held by external third parties, or in the case of a group decision that does not involve outward commitments, the commitment at that level is held by the group as a single entity. Social groups, both formal and informal, tend to be sensitive to these matters and act so as to discourage members from acting as if the group were not there as an intermediary. Outsiders are not to meddle with internal affairs, even if they have a stake in their resolution, except by addressing themselves to the group as a single entity – from the outside, as it were. Likewise, group members are not to act as their own agents where group-level commitments are involved. If a group fails to keep its commitment, but a single member did his part successfully, his defense is to be made in the context of the group and not directly to the third party. Only his group has the standing to accept his excuse as valid.

One kind of case stands out as a possible exception, though I think that in fact it can be most naturally explained on the two-level model I am arguing for. This is the case of a “whistleblower”: a member of a corrupt organization (often an employee of a corporation or government office) who passes information about illegal or unethical activities to legal authorities, journalists, or other outsiders. The proper way to think about the whistleblower, I think, is that she is not held to share responsibility for the group’s actions precisely because her own actions set her outside the group. She is, in effect, acting for the outsiders as an agent who has infiltrated the group structure. While her whistleblowing activities are generally admirable, it does not make sense to think of her as a member of the group in good standing. Recent history confirms this. American whistleblower Edward Snowden, for example, has been called “a traitor” by government officials of both parties,⁷ while Snowden himself claims he was acting on behalf of “the world I love” (Greenwall, MacAskill, & Poitras 2013): just the kind of shift in identification that the two-level model would lead us to expect.

It is generally understood that the whistleblower, in virtue of her subversive activities, is no longer a party to the group-level commitments, or at least not all of them. She may in fact remain devoted to the group’s ends, and continue to identify with the group itself, but strongly oppose the means which the group has adopted in pursuit of those ends. If she affirms this after the fact, or if her behavior otherwise suggests that she wishes to retain her group membership, she risks an especially sharp rebuke from within the group. The rebuke is justified, to an extent, because her actions can be seen as involving an implicit claim that she can remain a member of the group in good standing while abiding by some but not all of the group-level commitments.⁸ They will say that if she disagreed about means but not about ends, she should have raised those issues internally rather than getting outsiders involved.⁹ Such a rebuke would effectively be appealing to a two-level structure of commitments such as I have been describing, albeit not under that name. U.S. Secretary of State John Kerry, for instance, had the following to say about Snowden after Snowden expressed a desire to return to the United States: “A patriot would not run away and look for refuge in Russia or Cuba or some other country. A patriot would stand up in the United States and make his case to the American people” (Serrano 2014).

The fact that her actions open the whistleblower to such a rebuke does not preclude the possibility that she may still be doing the right thing, all things considered. But it is one of those unfortunate circumstances in which doing the right thing comes at a normative cost: she may have to violate the norms of group membership in order to conform with the norms of honesty, civic responsibility, or morality more generally.

I now want to return to my original cases of you, me, Amy, and Chad sitting on the couch. I believe that the two-level view of group commitment that I have been presenting can offer an illuminating explanation for the two ways in which the commitments in that group fail to be met. Recall that in Case 3, we all agree to go to the store together, and then no one moves when the time comes. We are now in a position to recognise this as a failure regarding a group-level commitment: a commitment the group has to itself as a single body. This diagnosis also gives content to the analogy, mentioned earlier, between Case 3 and Case 2, where I am

**4.
The
“whistleblower”
case**

**5.
Group and
personal failures
revisited**

⁷ For a small sample of reactions along these lines, see Weisman (2013) and Serrano (2014).

⁸ More likely, a rebuke of this kind will fit a pattern identified by Gilbert: someone has the standing to issue a rebuke but, all things considered, it is not morally appropriate to do so. See for example Gilbert (1993, p. 702) and Gilbert (2008, p. 7).

⁹ Of course, she may have already tried to do this, or it may for other reasons have been unfeasible. There are a number of delicate issues here that invite further exploration.

alone on the couch and fail to go at the time when I had previously decided to go. In both these cases the entity which formed the commitment (and later failed to act on it) was also the holder of it.

The disanalogy between Cases 2 and 3 is that in Case 3 there are also individual-level commitments which each of us owe to the others, to each do our part in pursuing the joint goal of going to the store together. On the one hand, since our individual roles were never clearly determined, none of us is in a position to rebuke anyone else in particular, since there is no determinate fact about who individually failed how or to what extent. On the other hand, each of us as a group member necessarily has at least the minimal commitment to act *in some way* so as to help bring it about that our group-level commitment is fulfilled. Given that none of us did *anything*, it is clear that each of us failed more or less equally in that individual commitment. So while it may not be appropriate to say “I failed you all” or “You all failed me”, it would be appropriate to say either “We failed each other” (individual-level) or “We failed ourselves” (group-level). But the most salient of these is the group level: “We failed ourselves”, or simply, “We failed”.

Case 4 is the reverse of Case 3 in that the most salient failure is on the individual level. As far as the group-level commitment, it is unclear to what extent we ought to say that *we* failed. As mentioned earlier, the original goal of all four of us going was not achieved. Still, three of the four did go to the store, and they did go together. Perhaps the three of you took my inaction as a signal that I was no longer invested in the group commitment, and by leaving without me you signaled in response – as you all had a right to – that you no longer regarded me as part of the group that was committed to going to the store together.¹⁰ In that case, any one of the three of you could say “We went to the store”, but I could not: I am no longer a part of that particular “we”. This in itself is not problematic, since groups regularly have members come and go without altering the identity of the group. On the present account this general phenomenon would be explained, at least in part, in terms of group-level commitments remaining unchanged while individual-level commitments are more responsive to the shifting boundaries of membership.¹¹

In any case, the important thing is that I have failed the three of you by not going. The two-level view confirms and explains this: however it works out at the group level, I have broken my individual-level commitments. I broke my commitment to each of you individually by not acting as a committed group member should, and I broke my commitment to the group by not doing my part in pursuit of the collective goal. Remember that while these commitments have overlapping or identical content, they have different holders: one is owed to the group members, viewed as individuals, and one is owed to the group viewed as a single unit. You could therefore rebuke me either as yourself, or on behalf of the group; it would be equally appropriate for you to say “He failed each of us” (individual-level), and “He failed all of us” or simply “He failed us” (group-level). What these rebukes have in common is that, whatever their source, they are directed to me in my capacity as an individual group member and do not implicate the group itself in what is ultimately *my* failure.

Conclusion Social life involves commitments, and some of these commitments can only be met when we acknowledge that they are not *ours* in an individual sense. Instead they belong to a group, and

¹⁰ There are interesting questions here about the ways in which a group can adjust its goals and its membership. Regrettably, I will have to leave them for another time.

¹¹ This, too, seems to me an area worth examining in greater detail; however, to do so here would distract from the main argument of the present paper. A closer examination of the interplay between group- and individual-level commitments in a group setting only makes sense once we have reason to believe that there *are* such levels of commitment, which is what I take myself to be establishing here.

it is only to that group and its members that we can be committed on a personal level. I have argued that everyday social practices reflect a tacit but real recognition of this distinction. It is true that a two-level structure of commitments can sometimes limit our actions in uncomfortable ways. On the other hand, it is just this structure of commitments that allows one, in a not purely metaphorical sense, to be a part of something larger than oneself.

REFERENCES

- Bratman, M. (1999), "Shared intention and mutual obligation", in M. Bratman, *Faces of Intention*, Cambridge University Press, Cambridge, pp. 130-141;
- Brewer, T. (2003), "Two kinds of commitments (and two kinds of social groups)", *Philosophy and Phenomenological Research*, 66(3), pp. 554-583;
- Gilbert, M. (1989), *On Social Facts*, Routledge, London;
- . (1993), "Agreements, coercion, and obligation", *Ethics*, 103(4), pp. 679-706;
- . (1997), "Group wrongs and guilt feelings", *The Journal of Ethics*, 1(1), pp. 65-84;
- . (1999), "Obligation and joint commitment", *Utilitas*, 11(2), pp.143-163;
- . (2006), *A Theory of Political Obligation*, Oxford University Press, Oxford;
- . (2008), "Social convention revisited", *Topoi*, 27(1), pp. 5-16;
- Greenwald, G., MacAskill, E. & Poitras, L. (2013), "Edward Snowden: the whistleblower behind the NSA Surveillance Revelations", *The Guardian*, June 11, 2013;
- Roth, A. (2004), "Shared agency and contralateral commitments", *The Philosophical Review*, 113(3), pp. 359-410;
- Serrano, R. (2014), "Former NSA contractor Edward Snowden says he was a trained spy", *Los Angeles Times*, May 28, 2014;
- Weisman, J. (2013), "Boehner calls Snowden a traitor", *The New York Times*, June 11, 2013;
- Wilkins, B. (2002), "Joint commitments", *The Journal of Ethics*, 6(2), pp. 145-155.