

Collective Responsibilities of Random Collections:
Plural Practical Self-Knowledge Among Strangers

Hans Bernhard Schmid, University of Vienna (hans.bernhard.schmid@univie.ac.at)

Draft version; 9/4/2016

There are things people should do, and many of those are things they can only do jointly, together with others. A well-organized society usually has corporate agencies for these tasks. Since crime intervention in larger communities is too big a task for a single Sheriff, we have our Police Department. To fight fires that nobody can put out by him- or herself, there is the Fire Department, the Ambulance is in charge of medical emergencies, for poverty and deprivation there are the Social Services, for moral catastrophes in remote corners of the world we have Foreign Aid, and so on. These organizations are the agents to take care of those morally sensitive tasks which no one can perform alone.

We hold these organizations responsible if they fail. In recent philosophical research, it has come to be widely accepted that we are justified in thinking of such organizations as collective agents of their own, and thus as collective bearers of responsibility (cf. List/Pettit 2011). One of the arguments for collective or corporate responsibility is that the responsibility in question often cannot be straightforwardly reduced to the responsibility of the individual members of the organization. Though there are usually individual members to blame, blaming them for failing to do their job, for not performing in their roles, is different from blaming them for the act or omission in question. The Chief of the Fire Department is responsible for failing to re-organize his department, but you may not simply blame the boss for not putting out the fire in the way you would blame an individual bystander who fails to act with all the necessary means to intervene at her hands. Thus it makes sense to assume that organizations do have responsibilities of their own.

Living in a well-organized society with functioning corporate agents, it is easier for us, as individuals, to live up to our moral responsibilities. All we have to do, as individuals, is to live up to our individual responsibilities, and if the task is too big for us, individually, to dial the right number – the moral responsibilities for many of those actions which we can only perform jointly is on the broad (if metaphorical) shoulders of our corporate agents. Yet even in near ideal circumstances, this is not always enough. Even if the Police, Ambulance, Fire Department etc. are well organized, they cannot be everywhere. In some cases, we expect of people to organize themselves, and live up to the duties we have in the domain of joint action by teaming up and doing the right thing. Even of a random collection of perfect strangers, we sometimes normatively expect to act jointly, at least until the relevant organization appears on the scene.

What kind of responsibility is at stake here? As this seems to be somewhere between the individual and the corporate case, one way to approach this question is from the distinction between the responsibility of individuals and the responsibility of corporate agents. Is our duty to team up and act jointly similar to our individual responsibility to dial the right number and to pay our taxes, or is it more like the City Police's responsibility to do their job? The literature largely seems to favor the first line (I.). An obvious argument for this view is that a random collection of individual simply isn't the kind of entity that is a suitable target for blame: a random collection of strangers is no collective agent, especially where the random collection fails to team up. And yet, there is something about many such cases in virtue of which there is not simply individual responsibility to *team up* at stake, but rather our responsibility *to act*, too. If we fail in these cases, it is not simply a failure at teaming up – what we are responsible for is *not acting*, and this responsibility is collective. In this paper, I

will identify a condition under which responsibility of the collective sort is involved among random collections of strangers. Where this condition is met, random collection of individuals can be the kind of entity that is a suitable target for blame even where they failed to team up. The condition in question is the right sort of practical knowledge of what the strangers should be doing together (II.). I further argue that this kind of knowledge rather than organization is at the heart of collective agency (III.).

I.

Here is a version of the classical example around which much of the current literature on the topic revolves – the classical case is in Virginia Held’s 1970 paper by the title “Can a Random Collection of Individuals be Morally Responsible?” The case is a random collection of individuals in a sparsely occupied subway car that witnesses a bully beating up a defenseless person in such a way that it is obvious that the victim’s health is in danger. Assume further that – in contrast to Held’s example – that the bully is obviously much stronger than any one single individual in the random collection of passengers, even if that individual could count on the victim’s active support. It is obvious that if anyone of the passengers were to intervene alone in this situation, and none of the others were to join in, the helper’s health, would be in danger, too. However, if at least two of the witnesses would intervene in a minimally organized way, it seems obvious that the abuse could be stopped. Imagine now that no one takes the initiative, so that joint action comes about. The core question is: what is the responsibility involved in these cases? One way to find out our intuitive view is by asking ourselves whom we would blame for what.

Let us assume that under these circumstances, no one can be blamed for not intervening in the victim’s behalf alone. Quite uncontroversially, however, we would blame each of the passengers for failing to taking the initiative to gather a group together. But is there also a sense in which we would blame the random collection as such for not intervening, even though it is not an organized group?

Virginia Held argues that “we would hold the random collection morally responsible for its failure to act as a group” (477). This seems right, as what the random collection failed to do is *to act*, rather than just to team up. However, Held then goes on emphasizing the distributive nature of this responsibility: “if random collection R is morally responsible for the failure to do A, then every member of R is morally responsible for the failure to do A, although, perhaps, in significantly different proportions” (ibid.).

This view seems problematic for two reasons.

1. The first reason has to do with the difficulty of individuating action A. If A is adequately described as something like “coming to the victim’s aid” (a), saying that the collection is responsible would amount to holding each individual responsible for something he or she is not morally required to do. This seems counterintuitive. If A is adequately described as “subduing the bully” (b), we would hold each individual responsible for something that individual is not even able to do, which is even more counterintuitive. Is A therefore simply something like “acting so as to transform the collection into an organized group” (c)? Held argues herself that “it may well follow that in some cases all the individuals in a random collection are responsible for not acting to transform the collection into an organized group, even though none of these individuals is responsible for not taking the action that ought to have been taken by an organized group in these circumstances” (480). Yet “acting to transform the collection into an organized group” seems to be something for which individuals are responsible, and not the collection as such. It thus seems that no one the three ways to individuate action A is convincing. a) and b) are too strong, c) is too weak. (a) blames

agents for having failed to do something they are not morally required to do, (b) blames agents for having failed to do something they are unable to do, (c) seems reasonable in and by itself, but it does not give us an account of the responsibility we're trying to explain; it is not "holding the random collection responsible for its failure to act as a group".

2. For the second objection against Held's distributive reading, imagine the following version of the subway scenario. You are one of the passengers. You immediately call the police. But you know that you have to do more. Given the strength of the bully, you're not morally required to start an intervention by addressing the bully, hoping for another witness to follow suit. Rather, what you should do is to turn to your fellow travellers, establish joint readiness and agree on intervening together – perhaps by saying "Let's stop this, shall we?", and wait for a nod from your prospective partner – and then do it, together with your partner. Imagine that for some reason, you just can't find any partner. You look around you. The other people simply avoid making eye contact with you. You call on them, perhaps by saying "we must do something – will anyone help me to stop this, please?", but to no avail. You do whatever you can to team up with some other witness to stop this, without putting your life in danger. But you fail. You just can't find any partner for an intervention. At the next subway station, the Police finally come and subdue the bully – but at that time, the victim has suffered further harm that could have been prevented by some decisive team action. Certainly, this is good reason for blame.

The point of this version of the example is that it is still true, in this case, that the random collection failed to act as a group in this scenario. But how can we make sense of the intuition that "we would hold the random collection morally responsible for its failure to act as a group" here? It is not the case that the responsibility can be distributed to *all* the individual members: after all, you did whatever you could to transform the collection into an organized group. It is not that you just have a lesser share in the distribution of individual blame. Rather, it seems that you're not to be blamed at all. This seems to leave us with three options. Either the collection is not responsible (i), or the collection is responsible, but you're not part of it (ii), or the collection is responsible, but not in the distributive way sketched by Held (iii).

Let's start with (ii). In our scenario, it seems natural for you to say that "they", the other witnesses, are morally responsible. Yet what this collection of individuals is responsible for is clearly not "its failure to act as a group". It's not the collection of the witnesses *minus you* that failed to act as a group, but rather the collection of the witnesses *including you* that failed to act as a group. The failing collection includes you, too. It is not the case that the collection should have acted jointly without you. Rather, the members should have acted jointly with you. For this reason, holding the collection responsible, but excluding you from the collection, does not seem plausible.

This leaves us with the option of letting go of the idea of the collection's responsibility (i), or assuming that the collection's responsibility is of the collective rather than the distributive sort (ii). Rather than arguing against the former alternative, I'd like to focus on the latter option. The problem is that in the core sense of the word at least, responsibility is for action. Where we hold people responsible for events or states of affairs, we do so in virtue of some wrongdoing (including omissions). In blame, we are addressing agents. Furthermore, we hold them responsible for *their own* actions. Random collections, however, do not seem to qualify as agents. None of the criteria listed in the recent literature on group agency seem to apply to random collections. There is no joint commitment of the sort Margaret Gilbert (1989) assumes in her account, there is, it seems, no "rational unification" of a single perspective required by Carol Rovane (2004), and no established decision procedure of the sort analyzed by Philip Pettit and Christian List (2011). Therefore it seems that our random collection in the subway car simply is not a suitable target of blame, because it is no (group, or plural) agent.

One way to go from here could be simply to detach collective responsibility from group agency. Sara Rachel Chant has made this suggestion in a recent paper (Chant 2015). Her fictional case is what she now calls a Hollywood Standoff – previously known under the politically a bit less correct label “Mexican Standoff”. This is a situation in which three agents hold each other at gunpoint in such a way that A threatens B, B threatens C, and C threatens A. The first to shoot will surely be shot, and if anyone lowers his gun, chaos would ensue, and anyone could be shot. Chant, in her scenario, has a child dying of thirst within sight. The three Hollywood standoffers could save the child by each one giving the child the last sip of water from his flask. In this situation, Chant argues, we have a pure case of collective responsibility without distributive personal responsibilities – no one in the collection of three is responsible for not lowering his gun and coming to the child’s help, let alone for the death of the child. But as a collection, they are still responsible for the child’s death. Chant argues that this thought experiment shows that there is collective, non-distributive responsibility without collective agency, as the three parties in the standoff are not a collective agent.

Chant’s thought experiment is striking in that it clearly exempts all participant individuals from individual responsibility. However, the central question concerning collective responsibility is basically the same in Chant’s and Held’s scenarios, and it does not seem to be addressed in Chant’s short paper: *In virtue of what* is the collection’s failure to do the right thing *the collection’s own omission*? The intuitive answer seems obvious: it is because these three people choose to threaten each other rather than to come to the child’s help. Perhaps we can attribute responsibility to entities that are not of the sort analyzed in current philosophical research on group agency; but surely, if we understand responsibility as related to praise and blame, there has to be a way in which the failure in question is the collection’s own doing. If we want to hold on to the intuition that such collections are collectively responsible, we need an account of how the doing in question is the collection’s own doing.

II.

There is indeed such a feature, but in order to see it, we need to dig a bit deeper into the structure of agency than most current accounts of group agency. Virginia Held shows where to start. In her discussion of the general notion of responsibility, Held points out that responsibility implies some sort of practical knowledge, knowledge of what it is the agent is doing:

“To hold an individual responsible for an act requires that he be aware of the nature of the action, in the sense that he is not doing A in the belief that he is doing B” (472).

Of course, this does not mean that you’re off the hook as long as you did not intend your action under the description under which it is morally reprehensible. It just precludes holding an agent responsible for what Held calls “thoroughly unascertainable” aspects of your action. To use Held’s example: if you throw an explosive through the window of a house, you’re morally responsible for the death of the inhabitant, even if you didn’t know that somebody was home. But if you ring a doorbell that unbeknownst to you has been connected to an explosive in the house, you’re not morally responsible for the death of the inhabitants. Held stops her discussion there, and simply examines how this translates to the case of a random collection. Arguing for the view that random collections of individuals can be morally responsible, Held argues such collections can indeed be said to have “awareness of the moral

nature of the action". This is an interesting suggestion well worth exploring; however, she hastens to add the following:

"If we say that, in special circumstances, a random collection can be aware of the moral nature of an act, we do not claim the existence of an inexplicable group awareness over and above the awareness of its individual members, only that we are sometimes entitled to say 'Random collection F is aware that p,' even though we cannot carry out a reductionist demand for statements about each individual member." (476)

The context of this passage suggests that Held thinks the non-reducibility of such statements is just due to pragmatic reasons only: we just can't say who's actually individually aware and who's not, that's all. No special sort of "group awareness" is needed, just some individual awareness.

Held here misses the crucial point concerning the kind of collective responsibility that can apply to random collections of individuals. To see the point, let us look again at the individual version of the knowledge or awareness in question. Remember that the knowledge or awareness in question is the feature in virtue of which an agent is responsible for an action. As we have argued, this feature must also be the feature in virtue of which an action is a subject's *own* action (in the sense in which it is only for their own actions that we blame agents). The "ownness" of the action in question has to be part of what is known – or of which there is awareness – by the agent, in the knowledge or awareness in question. In other words, this knowledge or awareness (or consciousness, in Held's sense) has to be *self-knowledge*, *self-awareness*, or *self-consciousness*, and it has to be self-knowledge, self-awareness, or self-consciousness of a special kind.

Knowledge, awareness or consciousness of ourselves can be *de re* or *de se*, observational vs. non-observational, reflective vs. pre-reflective, first-personal vs. third-personal (I take all of these terms and distinctions to be equivalent in the following). The *locus classicus* for a case of self-knowledge of the *de re*, observational, reflective, third-personal kind is Ernst Mach (1886, 3) who enters a bus and sees his own image reflected in a mirror, and without knowing that it is himself he observes, judges that his clothing is shabby. Looking at the mirror, he learns something about himself, but the knowledge in question is not of the first-personal kind, and it is thus deficient. The same is true for practical knowledge, and it pertains to the question of responsibility. Consider the following moral equivalent of Ernst Mach's case. Imagine the case in which Mach, again, enters the bus and sees himself mirrored in the opposite window without recognizing that it is himself that he sees. Now he observes that the man in the mirror is actually blocking an elderly person's way rather rudely and correctly judges that the man in the window is doing something wrong. He has clear knowledge of the moral nature of what he is doing, but he lacks the right kind of self-knowledge; he does not know that it is he who is doing the blocking. Even though there is knowledge of the moral nature of the act *of sorts*, this kind of knowledge does not constitute moral responsibility. The example illustrates that for somebody to be morally responsible for what she does or fails to do, it is not enough for that agent to know the moral nature of what she is doing or not doing *de re*, observationally, reflectively, or third-personally; she need to know the moral nature of what *she* is doing or not doing *de se*, first-personally, pre-reflectively, or non-observationally. All the knowledge of the moral nature of an action would not make somebody a *responsible* agent were it not for that kind of self-knowledge, self-consciousness, or self-awareness.

With this in mind, let us now return to joint action. Here is the central claim. In order to be collectively responsible for an action we, together, need to have awareness, or knowledge, of our action *as ours*. This knowledge, awareness, or consciousness is first personal, but it is first

personal *in the plural*. It is not a combination of each participant's self-awareness of what he or she is doing, individually, with some observational and/or inferential knowledge of what others are up to. I have argued in a different paper that while group self-knowledge, or plural self-awareness, certainly differs from individual self-knowledge, or singular self-awareness, it is of the same kind as the latter (Schmid 2014a). Self-awareness, self-consciousness, first-personal knowledge comes in two forms: singular and plural. Held ignores this because firstly, she does not account for the first-personal nature of the awareness or knowledge she requires for responsibility in general, and secondly, because she thinks that if that kind of awareness were different from individual awareness or knowledge, it would have to be "an inexplicable group awareness over and above the awareness of its individual members". The mistake to be found here, as in so many other cases, is the idea that if a group has awareness or knowledge, it has to be a collective singular awareness or knowledge. Yet the plural is not a collective singular. Plural self-awareness is something individuals have, not somebody else over and above their heads, only that in contrast to singular self-awareness, they have it only *together*, as a group.

III.

Let us now turn back to our subway car example to see how this spells out in terms of responsibility. Imagine we are that random collection of strangers. Take the asymmetrical scenario. You are the one member in the collection that does everything that is possible "to transform the collection into an organized group". I am one of those members who looks away and does not do anything. We have argued that the question of whether or not we are collectively responsible for the failure to help depends on whether this failure can be attributed to us, collectively, as *our* failure. And I have claimed that this depends on whether or not we, together, are plurally self-aware of the failure *as ours*. How do we find out whether or not this is the case? One condition of the plural awareness in question is joint attention. If you're aware of what is going on, and of the fact that assistance is needed, but you think that I'm unaware of it, your responsibility is to establish joint attention and readiness for action. If we're jointly aware of what's going on, and I refuse to participate, I am personally responsible for my refusal, and thus for our failure. Yet the failure is still *ours*. *We* failed, and are thus the proximate target for blame. But in this scenario, you're personally off the hook, because you did not do anything wrong – it is me who did wrong, and the wrong is of a special kind: I prevented us from doing the right thing. We're collectively responsible, but our collective responsibility is transferred to me personally. Thus I'm doubly guilty; I did not engage in mobilizing the group, and I prevented *us* from doing what is right. In Held's initial scenario, in which nobody engages in transforming the collection into an organized group, the structure of responsibility is different. Each of us is personally responsible for not taking action to mobilize the group. But it assuming joint attention, is also the case that we, together, are collectively responsible for failing, as a group.

Let me conclude with a brief remark on the connection between collective responsibility and group agency. It is obvious that a random collection of individuals is not a group agent of the received kind. There is no joint commitment, rationally unified perspective, or decision rule in this case. But even un-organized random collections can be plurally self-aware of what should be done as *that collection's own* actions of the sort that constitutes responsibility. This is the case wherever we, jointly, know first person plurally what we should or should not do, together. As I have argued elsewhere, plural pre-reflective self-awareness, or groundless group self-knowledge, is not in itself the kind of organization required for the received group agents. But it is the feature in virtue of which there is normative pressure towards rational

unification and organization among us. It is in virtue of our self-awareness that we are the owners of our attitudes and actions, in virtue of which there is pressure for formal and material unity of our attitudes, and in virtue of which our attitudes are our commitments (Schmid 2014, 2016). This has a singular or individual as well as a plural or collective form. And just as we are individually responsible for our actions, as agents, even if we fail to achieve full unity, we are collectively responsible even if we fail to organize ourselves. Plural agents do not exist in virtue of some achieved rational unification and institutionally established organization, but in virtue of the feature which provides the normative pressure towards rational unification and organization. In this sense, even random collections of individuals can, under condition of joint attention and a shared moral outlook, be collective agents, and thus indeed be collectively responsible.

References

- Chant, Sara Rachel (2015): "Collective Responsibility in a Hollywood Standoff." *Thought* 4/2, pp. 83-92.
- Gilbert, Margaret (1989): *On Social Facts*. Princeton, Princeton University Press.
- Held, Virginia (1970): "Can a Random Collection of Individuals be Morally Responsible?" *The Journal of Philosophy* 67/14, 471-481.
- List, Christian/Pettit, Philip (2011): *Group Agency. The Possibility, Design, and Status of Corporate Agents*. Oxford, OUP.
- Mach, Ernst (1886): *Analyse der Empfindungen*. Jena, Gustav Fischer Verlag.
- Rovane, Carol (2004): "What Is an Agent". *Synthese* 140, 181-198.
- Schmid, Hans Bernhard (2014): "Plural Self-Awareness". *Phenomenology and the Cognitive Sciences* 13, 7-24.
- Schmid, Hans Bernhard (2014a): "Expressing Group Attitudes. On First Person Plural Authority." *Erkenntnis* 79, pp. 1685-1701.
- Schmid, Hans Bernhard (2016): "Groundless Group Self-Knowledge." In M. Fricker/M. Brady (eds.), *The Epistemic Life of Groups*. Oxford, OUP, chap. 3.