

Groups with minds of their own making

Leo Townsend, University of Vienna

Prepublication version: please cite published version

Introduction

It was not very long ago that the idea of group agency—the idea that social groups might be agents in their rights, with minds of their own—was utterly repudiated by philosophers in the analytic tradition. Not only was this idea dismissed as manifestly false, it was also, on account of its perceived metaphysical extravagance, seen as intellectually backward or anti-scientific. And not only was it roundly dismissed as false and somewhat backward, it was widely viewed as a politically dangerous doctrine, one that was apt, if allowed to infiltrate the broader culture, to promote anti-democratic or totalitarian ideologies.¹

In recent times, however, the theory of group agency has been resuscitated and repackaged, to the extent that it is now a perfectly respectable view held by several respected analytic philosophers.² Perhaps the leading such figure is Philip Pettit,³ who holds that suitably organized groups do indeed qualify as agents, and that there is nothing metaphysically untoward about their doing so.⁴ And this is just as well, claims Pettit, because from a political standpoint, the recognition of groups as agents is integral to the proper functioning of a democratic society.⁵

In fact, Pettit goes even further that this, claiming that groups with the right structures may qualify not simply as agents but as full-fledged *persons*. This is because such groups not only possess “minds of their own” (Pettit 2003), but can also, as Pettit now puts it, “make up their minds” (Pettit 2014; 2016a). According to Pettit, it is this ability that allows group agents to take responsibility for what they think and do, and thus perform as *conversable* subjects or persons, ‘capable of being reached and engaged in speech’ (Pettit 2014, 112).

¹ The repudiation of group agency along these three lines is epitomised in the work of mid-19th century methodological individualists (see, e.g., Popper 1957; Watkins 1957).

² See, e.g., Gilbert (1989; 2014), List (2014; 2016), Tuomela (2013), Tollefsen (2002; 2015), Rovane (1998; 2014), Schmid (2014; 2016).

³ Note that the account I attribute to Pettit in this paper has been developed, in significant ways, through co-authored publications (notably McGeer & Pettit 2002; List & Pettit 2002, 2011; Pettit & Schweikard 2006). I do not wish to undervalue the contributions of Pettit’s collaborators, but given that he is the common thread in all the work I discuss here, and that I focus in large part on his most recent, single-authored work, I shall henceforth call the account under discussion simply ‘Pettit’s account’.

⁴ See Pettit (2003; 2007; 2014a, b), List & Pettit (2002; 2011).

⁵ See Pettit (2001; 2012).

Although Pettit's claim that groups can have "minds of their own" has attracted plenty of critical attention, his grounds for extending the laurel of personhood to groups have not been as closely scrutinised. In this paper I make these latter arguments my critical focus. Since Pettit has developed his case for group personhood in subtly different ways across a number of different publications, my first step is to carefully reconstruct his account, explaining first how he understands the two capacities he considers central to personhood—the capacity to "make up one's mind" (1.1), and the capacity to "speak for oneself" (1.2)—before showing how he thinks these can be manifested in groups (1.3). With Pettit's account duly reconstructed, I then turn to criticism, arguing that Pettit's construal of making up one's mind does not do proper justice to our first-personal self-understanding (2.1), nor to our characteristic interpersonal forms of engagement (2.2). This leads me, finally, to consider an alternative construal of "making up one's mind" (3.1) and "speaking for oneself" (3.2) that is associated with the work of Richard Moran and which, I argue, could usefully be extended to groups (3.3).

1. Pettit's account

1.1 "Making up one's mind" as self-regulation

In order to properly appreciate Pettit's account of group personhood, we should first clarify his approach to the concept of personhood itself. Here Pettit explicitly favours a "performative" conception over what he calls an "intrinsicist" conception (List & Pettit 2011, Ch. 8). The idea, put simply, is that what makes something a person is how it is able to perform, rather than its particular material (or immaterial) constitution. It is worth noting, however, that the question of which certain performances "make something" a person can be understood in two distinct senses. On the one hand, there is the question of which external performances *qualify* the creature as a person, while, on the other hand, there is the question of what inner functioning or capacities will *enable* it to perform in the requisite ways – a question, as it were, about the *makings* of the person. Though Pettit himself is not explicit on this distinction, I think he addresses both of these questions, and hence provides an account of both the external and the internal dimensions of personhood. The external dimension of personhood (what makes something *qualify* as a person) he identifies as the ability to "speak for oneself", while the internal dimension of personhood (what *enables* something to perform in those ways) is the capacity for "making up one's mind".

The notion of "making up one's mind" has often been invoked to capture notions of agency and responsibility in mental life—the idea that one's attitudes are, in some sense and to some extent,

*up to one.*⁶ Pettit cashes out this basic idea in terms of *control*. On his view, “making up one’s mind” means being able to exercise some control over how one’s beliefs and other attitudes are formed and revised. This is not a matter of picking and choosing one’s attitudes as one pleases, for that would imply an unwelcome and implausible voluntarism about attitudes like beliefs and desires. Rather, the sort of control that Pettit describes is a matter of being able to supervise and in various ways *regulate* the processes through which one’s attitudes are formed or revised. Although Pettit has described such “self-regulation” in slightly different ways over the years, I think there are three main ways in which he holds that people are able to exercise regulatory control over their attitudes. These are: by *reasoning*, by *research*, and by *intentional self-control*.

The form of self-regulation to which Pettit has devoted the most attention is the capacity to affect one’s own attitudes through higher-order reflection or “reasoning”. To be able to reason, on Pettit’s account, is to be able to reflect upon the propositional contents of one’s first-order attitudes, and form “metapropositional attitudes” in response. This is something that language-using creatures can do, because the sentences of a language provide a way of “objectifying” the propositional contents of thought, thus allowing the reasoner to focus on the distinctively propositional properties of her attitudes in an effort to make them better satisfy standards of rationality (standards such as evidential responsiveness and consistency). For instance, she can consider whether some proposition she now believes is well supported by evidence, and if she finds it is not, then she may be led to abandon the belief. Or, she may wonder whether some subset of her beliefs is consistent, and, finding it not to be, she might be prompted to give up one (or more) of these beliefs. In this way, the activity of reasoning gives the agent a degree of rational control over her attitudes.⁷

Closely connected with the ability to reason is the ability to conduct inquiry or *research*, that is, to actively seek out evidence relevant to one’s actual or would-be attitudes. Say, for instance, that having recently moved to Austria, I wish to be better informed about the country’s politics. This is something I can accomplish by means of research: I can produce beliefs about Austrian politics in myself by reading the newspaper, talking to colleagues, watching the news, and so on. And I can do this more or less carefully, to the betterment or detriment of the rationality of the beliefs I thereby produce. If I exercise no care, for example by simply believing everything I am

⁶ See, e.g., Bilgrami (2006), Hieronymi (2006; 2009), McGeer (1996; 2008), Boyle (2011), and of course Moran (1999; 2001; 2012), whose work is discussed later.

⁷ McGeer & Pettit (2002), Pettit (2007)

told about Austrian politics by one outspoken neighbour, then my beliefs may not end up being well supported by evidence.

One final way, according to Pettit, in which people exercise regulatory control over their attitudes is by means of *intentional self-control*.⁸ The need for this form of self-regulation typically arises when the other two forms of regulation, reasoning and research, fail to produce their proper effects. Consider someone who, on the basis of painstaking research and reasoning, concludes that eating meat is wrong, yet finds it hard to maintain the belief that meat-eating is wrong. This is something she could address by resorting to techniques of self-control. She might, for instance, make a mental effort to remind herself of the solid arguments that eating meat is wrong, or she might start shopping only in meat-free stores, or she might avoid sharing meals with her carnivorous friends. This doesn't guarantee that she will retain her belief that meat-eating is wrong, but it does make it easier for her to do so, and so amounts to a sort of indirect control over her attitudes.

In recent work Pettit has claimed that being able to make up one's mind in these ways gives people a special form of self-knowledge that he calls "maker's knowledge" (in contrast with "observer's knowledge"). The basic idea is that someone who can make up her mind is in a position to know her own attitudes without resorting to introspective self-observation or behavioural self-analysis. When asked whether she believes that *p*, she needn't attempt to *discover* her attitude, but can instead *determine* that attitude simply by considering the matter of *p* itself. In this way, Pettit claims, the capacity to make up one's mind places people in a "particularly secure" (Pettit 2016a, 14) position to know their own states of mind. And this "maker's knowledge" plays a crucial role in supporting the discursive practices that will be described in the next section—the way that people lay claim to, and are generally credited with, the authority to "speak for themselves".

1.2 "Speaking for oneself" as self-representation

Just as the notion of "making up one's mind" is commonly connected with a sort of *intrapersonal* authority, so the notion of "speaking for oneself" is often associated with a sort of *interpersonal* authority. Pettit's approach here is heavily influenced by Hobbes, in that it associates "speaking for oneself" with a certain kind of authoritative *self-representation*. Hobbes held that the hallmark of personhood was the ability to give one's word to others—that is, to

⁸ McGeer & Pettit (2002), Pettit (2016a,b)

express one's attitudes to others in a way that licenses them to hold one responsible if one fails to live up to the attitudes expressed. To do this, according to Hobbes, is to "personate" oneself, or to act as one's own *spokesperson*.⁹

Pettit elaborates this idea by distinguishing between three basic ways of communicating one's attitudes: by reporting, by avowing, and by pledging. To report one's attitude is simply to say of oneself that one is so minded. This is something you might do, for instance, in a therapeutic context, when you find that you harbour certain unwelcome desires, or that you hold certain unreasonable beliefs. Importantly, when you merely report your attitude you do not identify yourself with it, and hence you limit your accountability. Should you turn out not to have the attitude ascribed, you would not normally face any sanction or rebuke, but even if you did, you could legitimately try to get yourself off the hook by claiming either that the psychological evidence you used to adduce the belief was misleading (the "misleading mind excuse"), or that, while the report was true when made, things have subsequently changed (the "changed mind excuse").¹⁰

By contrast, to avow or pledge one's attitude is not only to self-ascribe the attitude but furthermore to *commit* oneself to having the attitude, to "bet on oneself" (Pettit 2016b, 227) to display it. To see what is involved in such commitment, consider how Pettit understands the speech acts of assertion and promising. When you assert that *p*, according to Pettit, you therein avow the belief that *p*. This means that, should you fail to maintain and manifest the belief that *p*, you render yourself liable to rebuke or sanction from your interlocutor. And in the face of such rebuke or sanction, you will no longer have recourse to one of the excuses that would have been available had you merely reported your attitude. Specifically, you cannot appeal to a "misleading mind", claiming that it had seemed to you that you had the belief that *p*, but that appearance, it transpires, was misleading (although you might plausibly claim to have changed your mind). And promising, which involves pledging an intention, raises the stakes even further, since it involves foreclosing both of the excuses mentioned above. If you fail to live up to a promise, you can neither appeal to misleading psychological evidence concerning the intention pledged, nor claim to have changed your mind in the interim. Hence, your commitment in the case of a pledged attitude is not just to having the attitude but to *continuing* to have it.

⁹ See Pettit (2009, Ch. 4).

¹⁰ See Pettit (2016a, 15-16); Pettit (2016b, 222-227)

The fact that pledges and avowals are, in this way, ‘costlier’ modes of attitude-expression is precisely what Pettit thinks makes them attractive for both speaker and audience. Given that the speaker manifestly exposes herself to more risk, the audience will tend to find her avowals and pledges more credible than reports; and the speaker, recognising this, will find them more useful than reports for inducing the audience’s reliance. Or at least this will be so just so long as there is a mutual understanding between them that the speaker stands in a special relation to the attitudes she expresses. She can safely and successfully renounce certain excuses for failing to have and maintain those attitudes only because whether she has and maintains them is, to a certain extent, *up to her*. If she didn’t stand in this relation — if she couldn’t make up her mind, and so didn’t have a “maker’s knowledge” of her attitudes — then the confidence she exhibits by “betting on herself” in avowal and pledging would simply be false bravado. It is in this way that the *intrapersonal* capacity to make up one’s mind provides crucial support to *interpersonal* practice of speaking for oneself, making it possible for people to commit to others and enter contracts with them through ordinary speech acts such as assertion and promising.¹¹

1.3 Self-regulating, self-representing groups

We have seen, then, that on Pettit’s account persons are those agents who can “speak for themselves” by avowing and pledging their attitudes, and this is something that is made possible by the ability they have to control their own attitudes or “make up their minds”. We can now examine how, according to Pettit, suitably organized social groups manifest both of these aspects of personhood: how they manage to be both self-regulating and self-representing.

Let us begin with the capacity groups have for self-regulation. We saw earlier that one central form of self-regulation was “reasoning”—the ability to affect one’s first order attitudes by reflecting on their propositional contents and forming “metapropositional attitudes” in response. Now Pettit argues that groups can reason in exactly this sense, by employing organizational mechanisms such as the ‘straw vote procedure’.¹²

Before we examine the straw vote procedure, it is worth mentioning that Pettit considers *voting* to be the paradigmatic way in which groups generate their attitudes. In ideal cases what happens is that the individual group members consider some group-relevant question, form a personal judgment, and then vote in line with their judgment. These votes are then processed (or

¹¹ Pettit (2016b, 235-8).

¹² Pettit (2007), List & Pettit (2011).

‘aggregated’) in some way in order to arrive at a ‘group-level judgment’, which can then be counted as the group’s attitude on the issue in question.

Though Pettit sees voting as the paradigmatic method of forming group attitudes, he is nonetheless emphatic that simple voting procedures will be inadequate for this task. The main reason for this is well illustrated by the ‘discursive dilemma’, which shows that straightforward, issue-by-issue voting on a set of logically inter-related propositions cannot guarantee a group view that satisfies basic standards of rationality, even when each voter is herself individually rational.¹³ In other words, individual rationality plus this sort of voting does not, on its own, secure collective rationality. A canonical version of the dilemma demonstrates this by indicating how a group of three people might vote on the propositions P, iff P then Q, and Q so that the majority-favoured set ends up being inconsistent without any one of the individual voters having supported an inconsistent set:

	P?	Iff P then Q ?	Q?
Member A	Yes	No	No
Member B	No	Yes	No
Member C	Yes	Yes	Yes
Majority	Yes	Yes	No

It is in order to guard against this sort of collective irrationality that Pettit thinks groups should go in for more sophisticated voting procedures such as the straw vote. When a group’s attitudes are determined via a straw vote the majority-supported views are recorded but not immediately ratified. Members are kept in the loop on the results of their voting, and this allows them to form metapositional attitudes about the group’s first-order attitudes as they form. This sort of constant surveillance and reflection on the group’s actual or would-be attitudes can help to avoid rational problems like inconsistency. For example, if our group in the above example finds that they have voted not-Q, having previously supported both P and iff P then Q, then the ability of their members to form metapositional attitudes about these propositions will likely lead them to the realisation that they cannot, on pain of irrationality, endorse all three propositions. Hence they may re-assess all three propositions and change the group view on at

¹³ List & Pettit (2002, 2011), Pettit (2003)

least one, depending on what the evidence requires. In this way the adoption of a straw vote procedure, or something like it, can allow a group to regulate and improve the overall rationality of their attitudes.

Let us now turn to the ‘external’ dimension of group personhood, the capacity of groups to ‘speak for themselves’. This is also something that Pettit thinks can be facilitated by organizational structures, in this case by authorizing a spokesperson to speak in the name of the group. So just as collective decision-making processes such as straw voting allow to group to become of one mind, so the public use of a spokesperson allows the group to speak with one voice.

We saw earlier that “speaking for oneself”, on Pettit’s view, was a matter of expressing one’s attitudes to others in costly, responsibility-increasing ways. According to Pettit, the role of spokespersons in many organized groups is to speak for the group in exactly this sense. Spokespersons publicly announce the group’s decisions, judgments and intentions, and it is recognised both within and without the group that these words function to bind or commit the group, that the attitudes expressed are ones the group can be expected to maintain and manifest. This is clear, Pettit argues, from the spirit in which spokesperson’s declarations are offered and received in common practice:

“Is it excessive to take the declarations of spokespersons to be avowals and promises? Absolutely not, for the declarations are taken in common usage to rule out excuses of misleading evidence or change of mind that is characteristic of avowals and promises. Suppose a group fails to live up to a belief or value ascribed by an authorised spokesperson. It will not do for that spokesperson to excuse what was said on the grounds of having mistaken the evidence about what the group held. The spokesperson’s only recourse will be to resign from the role assigned by the group or, maintaining that role, to try to offer another excuse for the failure or to make an apology on the group’s behalf” (Pettit 2014a, 111)

The use of a spokesperson thus enables groups to avow and pledge their attitudes. To the extent that they are prepared to take on the costs of these modes of speech, and to the extent that they are recognised by their would-be interlocutors as being in a position to take on those costs, they should therefore, according Pettit, be counted as conversable, responsible subjects or ‘persons’.

2. Two problems in Pettit's account

2.1 The problem with self-regulation

My aim so far has been purely exegetical. I have reconstructed Pettit's two-part model of personhood, and I have shown how it is, according to him, that groups can perform as this model requires. I now wish to raise two problems for Pettit's account, as I have reconstructed it. The first targets his construal of making up one's mind as self-regulation, while the second targets his conception of speaking for oneself as a matter of self-representation. Though both of these problems concern Pettit's model of personhood itself, rather than his claim that groups can fit this model, I think that reflection on the group case can help us to highlight what is wrong in the model.

Consider first Pettit's way of construing "making up one's mind" in terms of a capacity for self-regulation. Here, for example, is how Pettit illustrates the process of metapropositional reasoning, and the rational benefits it can confer:

"Suppose I currently believe that *p* and that *q*. Perhaps because of worrying about the conditions in which I formed those beliefs, I may ask myself whether 'p' and 'q' are consistent propositions, setting in train a process of forming a belief in answer to the question and adjusting to the result" (Pettit 2007, 499)

In a more recent paper, Pettit details a step-by-step method for making up your mind on some proposition (Pettit 2016a, 10-11, bullets in original):

- "Make sure that you are clear about the meaning of the sentence or proposition.
- Identify the available data relevant to the proposition, taking care to overlook nothing.
 - Defer to the data. Letting them determine whether you assent, dissent, or withhold judgment.
 - If you assent to the proposition, then you believe it, and if you dissent from the proposition, then you disbelieve it.
 - Or at least this is so when things well."

Passages such as these help to clarify just where Pettit locates the activity and passivity involved in making up one's mind. What is up to the agent herself is which propositions she considers,

and how carefully she does so. By considering whether certain propositions she already holds true are consistent, or by exposing herself to the data relevant to some as yet unconsidered proposition, she “sets in train” certain non-intentional processes that culminate, when things go well, in attitudinal changes. But as far as these processes and changes go, the agent herself figures only as a kind of bystander—one who needs to wait and see which attitudes are produced in her. This is where I think the problem with Pettit’s model of making up one’s mind lies.

To be clear, the problem I am raising here is not essentially concerned with the *degree* of agency but with the *kind* of agency involved in making up one one’s mind, and what this implies about the relation in which the agent stands to her attitudes. More specifically, I think that passages like those quoted above reveal that, for Pettit, the agent who makes up her mind does so by exerting a kind of ‘external control’ over her attitudes. That is, she tries to *affect* her attitudes through reasoning, research or intentional self-control, these being effective means of producing rationally welcome changes. Hence, the agent must be seen, and must see herself, as standing at a certain remove from her own attitudes, in order that she may affect them in this way. This is not, I think, the ordinary, intuitive conception of “making up one’s mind”—the conception that is meant to capture a sense of agency over one’s attitudes that is central to personhood. On the contrary, when one makes up one’s mind about something (in the ordinary, intuitive sense), it is *that thing* rather one’s *mind* (or states of mind) with which one is centrally concerned, and whatever changes thereby take place amongst one’s states of mind are simply the expression of one’s take on that thing.

This is not, of course, to deny that we do sometimes need to adopt exactly the kind of ‘external’ perspective on ourselves that Pettit describes. It is a familiar human predicament to find that sometimes the results of one’s deliberation fail to “take hold in one’s psychology”, that one’s considered judgments just do not become one’s actual attitudes. For instance, a hopeless gambler might judge, on the basis of careful consideration of relevant evidence, that the gambler’s fallacy is indeed a fallacy, and yet, when she stands at the roulette table, she finds that judgment doesn’t stick as her actual, action-guiding belief. On the contrary, she finds herself susceptible to the fallacy: it just seems clear to her that, for example, black *is* more likely after a long run of red. And someone who finds herself in this kind of situation, might, as Pettit suggests, need to make use of various techniques of self-regulation or self-manipulation. She might need to avoid those situations or stimuli that tend to dislodge her considered judgment; she might need to recruit her

friends into some kind of policing role; or she might even resort to hypo-therapy or something similar in order to get her belief to match her judgment.

To insist that making up one's mind be associated with the sort of case in which a person's deliberations issue directly in her attitudes is not to deny that, often enough, we do find ourselves in this kind of alienated situation with respect to our own attitudes. It is just to deny that what one is doing, when one resorts to the techniques of self-manipulation and self-regulation called for in such situations, deserves to be called "making up one's mind". On the contrary, it seems that the question of whether to exert such 'external control' over one's attitudes only arises when the regular rational authority a person has over her attitudes has somehow failed—in other words, when she finds herself, for one reason or another, unable to make up her mind.

This external perspective on one's own attitudes also permeates the "maker's knowledge" that Pettit thinks is facilitated by making up one's mind:

"you yourself do not have to scan [behavioural and contextual] evidence in order to ascertain whether you believe that p, and neither do you have to scan the evidence about your states of mind on the basis of an introspection [...] You just have to know that you have made up your mind that p on the basis of a careful review of data. That knowledge serves in combination with the knowledge that painstaking assent generally generates belief to give you a basis for knowing that, all going well, you believe that p" (Pettit 2016a, 13)

This makes it seem as though the reflectively-available fact that one has made up one's mind that p is something that might serve *for oneself* as evidence that one believes that p—that one might infer that one believes that p on the basis of knowing one has carefully given one's assent and that "painstaking assent generally generates belief". But this is radically out of step with the way we usually take ourselves to know our own attitudes. If I have made up my mind that p then there is no question, for me, of needing to muster psychological evidence, such as the fact that I have made up my mind and the purported general psychological connection between the terminus of that activity (judgment or assent) and belief, in support of this self-ascription. In fact, to the extent that I do need to resort to such evidence in order to figure out my state of mind it will seem to me that I *lack* control or authority over these attitudes. So if Pettit's idea of "maker's knowledge" involves, as it seems to, a person treating the fact of having made up her

mind as evidence in support of her self-ascriptions, then it is a peculiarly alienated form of self-knowledge.¹⁴

It is worth noting that, if the concerns I have raised about Pettit's conception of making up one's mind and the "maker's knowledge" that is supplied by doing so are compelling, then this poses particularly acute challenge to Pettit's account of corporate personhood. This is because the features of Pettit's model of making up one's mind I have objected to are precisely what makes the model so fitting for the kind of organized groups with which he is preoccupied. From the perspective of group members who must vote on some issue, the 'attitudes' of the group are indeed things that they can, at best, affect. Of course, these members may enjoy a profound influence over what these attitudes will be: beyond their voting privileges they may also get to decide which propositions are on the agenda, what decision-making protocol will be adopted, and so on. But no matter how sizeable their influence, the use of such organizational mechanisms means that these individuals invariably stand at a remove from the attitudes of the group. These attitudes are not a direct expression of their rational activity, but rather things that they, as participating members, can only monitor and manipulate. And this also means, of course, that whatever "maker's knowledge" of the group's attitudes is possessed by group members is not a direct, spontaneous product of their participation in the group's decision-making processes, but rather something that rests in turn on their prior knowledge about how those processes work. All of this captures neatly the perspective of a member of an organized group, but it does not, I contend, do justice to the first-person perspective of a single person who makes up her mind. So either Pettit needs to adjust his model of making up one's mind—in which case it may no longer be achievable by groups—or he needs to admit that the model does not do what the notion of making up one's mind is meant to do, that is, capture the intuitive sense in which our attitudes are up to us.

2.2 The problem with self-representation

¹⁴ I do not of course deny that sometimes we do come to know our attitudes on the basis of psychological evidence. We do not always immediately know what we think of something, and in such situations it can be that the best way to figure out what we do think is not by means of deliberation but by, for instance, considering our own past behavior, along with our present inclinations and dispositions. For example, under the promptings of a trained psychoanalyst I can learn that I believe I am responsible for my parents' divorce, even though, were I to deliberate on the issue of whether I am responsible for my parents' divorce, this is not the conclusion I would reach. The criticism I wish to raise against Pettit's account of maker's knowledge is not, then, that the phenomenon he describes under that banner—which ultimately involves coming to know one's mind on the basis of psychological evidence about oneself—is impossible. Rather, my complaint is that this is not the kind of self-knowledge that is essential to personhood, the kind that intuitively deserves the label of "maker's knowledge".

I turn now to a second problem with Pettit's account, concerning the way he construes "speaking for oneself" as a matter of self-representation. Pettit's idea was that to avow or pledge one's attitudes is to do more than simply report them—it is to commit oneself to them, and hence to undertake some kind of responsibility for them. The problem, however, is that the form of responsibility for our attitudes engendered by Pettit's account is not actually the fundamental form of responsibility our discursive practices presuppose.

To see this, consider what someone who avows her belief thereby makes herself responsible for. Say, for example, that Per tells you that Oslo is south of Bergen. It seems natural to say that, in doing this, Per certainly undertakes a form a responsibility: he makes himself answerable to you for the truth of the claim that Oslo is south of Bergen. This means that regardless of whether Per believes Oslo is south of Bergen, or whether he evinces this belief in his behavior, you acquire the standing to hold Per responsible, by questioning or challenging him on the truth of the claim he has made. And being answerable, Per is obliged to answer your challenges, not by giving an explanation or excuse as to why he does or does not have a particular belief, but by giving reasons in support of the claim itself, that Oslo is south of Bergen.

On Pettit's view, however, the commitment Per undertakes by "speaking for himself" in the mode of avowal is not this kind of *content-related* commitment, but is instead a kind of *attitude-related* commitment. According to Pettit, by avowing the belief that Oslo is south of Bergen, Per has not undertaken this 'content-related' form of responsibility, but has instead undertaken only an 'attitude-related' responsibility—a responsibility in respect of his believing Oslo is south of Bergen. Having avowed this belief, he is responsible for seeing to it that he has and maintains it. This is clear from the way Pettit analyses the normative element involved in avowal in terms of the 'excuses' that avowal forecloses for the speaker. For Pettit, the primary normative upshot of choosing to avow, rather than simply report, one's attitude is that this forecloses the 'misleading mind excuse': should it come to light that one does not have actually the attitude one has avowed, one cannot excuse oneself by claiming that one was mistaken about one's own attitude. But this excuse bears only on the question of whether a person actually has an attitude she represents herself as having. Hence, for Pettit, it is only with respect to your possession of the attitude that your decision to speak in the mode of avowal commits you, places you more firmly on the hook, making you, in that sense, 'more responsible'. The responsibility Pettit associates with avowal, then, is the responsibility a person has to match the persona she projects to others.

Now obviously people are often fit to be held responsible in this way, for failing to live up to the attitudes they have avowed. If Per tells you Oslo is south of Bergen but then tells others Oslo is north of Bergen, then you may reasonably demand an explanation. And if I warn you that chocolate kills, but you find me scoffing all the chocolates, then you may well be entitled to blame or criticise me for failing to practice what I preach. It may even be that there is *always* some element of attitude-related responsibility incurred when a speaker avows her attitudes. My concern with Pettit's account is not that I think there is no place for attitude-related responsibility in our discursive practice, only that we should not, as I think he does, neglect the crucial content-related form of discursive responsibility. The fact that we hold each responsible for the content of the attitudes we avow, and that acknowledge the propriety of epistemic challenges and demands for justification, is central to the way we use speech, and clearly our discursive practices would be unrecognizable without it.

Here too it is worth noting that this problematic feature of Pettit's account of "speaking for oneself"—the way it underpins only an attitude-related form of responsibility—might once again be what makes the account well-suited to organized groups such as corporations and political bodies. Unlike individuals, it seems plausible that the primary function of much group speech, or at least much of the speech that performed by spokespersons in the name of *certain* kinds of groups, is to set forth a kind of authorised portrait of that group's attitudes, i.e., where it stands on various important matters. Whereas individual people usually use speech to make discursive commitments to one another in respect of how things are in the world, corporations and political groups may, by contrast, tend to use speech primarily as a tool for self-representation. If this is right, then it becomes a central requirement—perhaps *the* central discursive requirement—for such groups that they live up to the persona their spokesperson has projected in order to foster the reliance of their public, their customers, their investors, or their stakeholders. But while this form of "speaking for oneself", and the attitude-related form of responsibility it supports, fits well with the discursive practices of certain kinds of groups, it does not seem to adequately capture the norms of discourse more generally. And so Pettit faces the same sort of dilemma as before: either alter the conception of speaking for oneself, in which case it may not end up so suitable for groups, or maintain a conception that fails to underpin the sort of authority and responsibility involved in regular discourse.

In this section I have argued that Pettit's way of thinking about "making up one's mind" and "speaking for oneself" does not do proper justice to the intrapersonal and the interpersonal

dimensions of personhood. This does not, of course, mean that the notions of “making up one’s mind” and “speaking for oneself” cannot be construed differently, so that they do properly capture what is distinctive of personhood. In the next section I describe an alternative construal of these two notions, associated with the work of Richard Moran, and consider whether it might be fruitfully extended to social groups.

3. Moran’s account

3.1 “Making up one’s mind” as self-constitution

One philosopher who has developed the notion of “making up one’s mind” in a very different way from Pettit is Richard Moran.¹⁵ The key difference between their accounts is that whereas Pettit spells out “making up one’s mind” in terms of *control*, Moran makes use of the notion of *authorial power*. Someone who makes up her mind, according to Moran, does not control, manipulate or regulate her states of mind; rather, in the course of her ongoing rational activity, she authors or constitutes them.

The central rational activity that Moran associates with making up one’s mind is deliberation. When a person deliberates, she forms or revises her attitudes by settling certain questions for herself on the basis of whatever considerations are relevant to the question. If she deliberates and settles for herself the question of whether *p* is true, she therein believes that *p*; if she deliberates about what to do and concludes that she should *phi*, she therein intends to *phi*.¹⁶ She therefore makes up her mind in the sense that these attitudes—her belief that *p*, her intention to *phi*—are the “precipitate of her ongoing rational activity” (Moran 2012: 219).

At first blush, this construal of “making up one’s mind” may not seem all that different from the ‘reasoning’ and ‘research’ described by Pettit. But there is in fact a major difference between these two ways of construing the process in question, which we can begin to appreciate by noting *how* each conceives of that process as resulting in attitudes (or changes in attitudes). For Pettit, the process of making up one’s mind through reasoning, research and intentional self-control is one that works only indirectly or mediately. Whatever steps a reasoner takes in order to regulate her attitudes, there is always a gap between her taking these steps and the production or revision of the attitudes in question. It is always possible, for instance, that she judges that *p*,

¹⁵ See esp., Moran (1999; 2001; 2012).

¹⁶ This is not to deny that sometimes we deliberate and find that what we cannot make our attitude match our judgement. It is simply to deny that in such a case one can be said to have actually settled the question about which one was deliberating, that, to have actually made up one’s mind.

perhaps on the basis of reasoning and research, and yet finds that the belief that *p* “may not gain a suitable hold in [her] psychology” (Pettit 2016a, 13).

By contrast, on Moran’s account, the process of deliberation issues directly and immediately in attitudes. When one considers the question of whether or not *p* and concludes that *p*, one therein believes that *p*. This is part of what it means, according to Moran, to conceive of oneself as a deliberator or rational agent:

“it means, for instance, that the conclusion of my practical reasoning shall count for me as telling what I will in fact do, or that the conclusion of my theoretical reasoning shall count for me as what I do in fact believe. This is in contrast to the situation of manipulating my own attitudes, where once I settle on what attitude to adopt I would then need to find a way to make this happen, finding the psychological means to implement my choice” (Moran 1999, 197).

Note that when Moran claims here that one’s judgments shall “count for one” as one’s actual attitudes, he is not saying that this is only how things will *seem* to the deliberator herself, from within the first person perspective. If that were what he were saying then his account would be fully compatible with Pettit’s account, since how things seem from within the first person perspective may not be how they actually are. Rather, the claim that the conclusions of one’s deliberations *count for one* as one’s attitudes serves to characterise the sort of power that the deliberator exercises in deliberating. It is not a power to make things happen—not a matter of control or regulation—but a power to take a normative stance, or “resolve oneself”. When one does this, one thereby authors the commitments constitutive of one’s states of mind:

“in a central range of cases, the person can be seen as the author of the state of mind itself, in the sense of being the person responsible for it. This is the person responsible for its justification and coherence with the rest of what she believes, and whose commitment to its truth makes the difference between this state of mind being one of belief, rather than a wish, a supposition, or a passing thought” (Moran 1999, 184)

Does this notion of authorial power over one’s attitudes—the idea that “the person can be seen as the author of state of mind itself”—imply any kind of implausible voluntarism about attitudes such as belief and desire? It should be clear that it does not. Voluntarism is the thesis that one

can believe (or desire) ‘at will’, that is, simply by deciding to, in much the same way that one can raise one’s arm simply by deciding to. But the agency that Moran claims is involved in making up one’s mind is not this kind of ‘productive agency’, but rather agency of an altogether different sort—what he calls “agency as responsiveness to reason” (Moran 2012: 135). For him, being an agent with respect to one’s states of minds does not mean being in a position to intentionally influence one’s attitudes, or being at liberty to produce them in oneself just as one likes. Rather, to be the author of one’s attitudes is to be such that the question of what one’s attitudes are can be, and often is, determined by what one has reason to believe, desire, intend, etc. It is in virtue of the fact that people usually take themselves, and are taken by others, to be standing in this relation to their attitudes that they are considered answerable for their attitudes, in the sense that they must be ready to justify them if challenged, draw suitable inferences from them, and ensure their overall coherence.

Apart from successfully avoiding voluntarism, Moran’s account is also not susceptible to the concerns I raised against Pettit’s account. Someone who is able to make up her mind in the sense of authoring her attitudes is not generally in the alienated situation of needing to hope that her judgment about what she should believe “takes hold” in her psychology. For Moran, in the ordinary case of making up one’s mind, the judgment a person reaches about what she should (i.e., has most reason to) believe simply *is* her belief on the matter, and there is no further thing she needs to do, beyond deliberating, to make it her actual belief. This is the authority we take ourselves and others to have as deliberators, and without which our doxastic and discursive lives would be unrecognizable. This is not to say, of course, that there are never cases in which one finds that one’s attitude does not match what one takes oneself to have most reason to think, and hence in which one may need to take some steps to somehow bring one’s attitude into line with one’s better judgment. But Moran’s point is that while such situations are clearly possible, indeed commonplace, they should not be seen as emblematic of the agency we take ourselves and others to have with respect to our attitudes. On the contrary, these are cases of alienation, cases in which one’s ordinary ability to make up one’s mind by deliberating has faltered, and it is a virtue of Moran’s approach that he is able to identify them as such.

Similarly, on Moran’s account, in contrast to Pettit’s, the knowledge an agent has of her own attitudes does not ordinarily rest on her awareness of the “general connection” between judgment and attitude. Instead, her “maker’s knowledge” is simply a form of *practical knowledge*: the knowledge a person has of what she is doing when she does it. In the standard case, I know

that I believe that *p* simply because I can deliberate and settle for myself the question of whether or not *p*. Such self-knowledge does not depend on my thinking that there is some reliable psychological connection between my judgments and my attitudes. On the contrary, if I could only know my attitudes by supplementing my deliberation with some psychological principle of this sort, I would not be taking myself to have the sort of authority Moran associates with making up one's mind. This is not to say, of course, that there are not cases in which one needs to figure out one's state of mind with recourse to psychological evidence. But the point is just that such cases, while possible, do not exemplify the special sort of self knowledge made possible by the rational authority associated with making up one's mind.

3.2 “Speaking for oneself” as illocutionary authority

We have seen some of the differences between Moran's conception of “making up one's mind” and Pettit's. Like Pettit, Moran also connects the intrapersonal capacity to “make up one's mind” with the interpersonal capacity to “speak for oneself” or (as Moran prefers to put it) to “speak with authority”. But Moran understands both the nature of this capacity and its connection to making up one's mind in a markedly different way from the way these are understood by Pettit. For Moran, it is not that making up one's mind *emboldens* one to speak with authority; rather, speaking with authority involves exercising the very same sort of authorial power that is involved in making up one's mind. When one speaks for oneself, one lays claim to, and is credited with, the *illocutionary authority* to constitute one's utterance—make it count—as a certain kind of speech act. And it is only because of the fact that people lay claim to and are credited with such authority as speakers that they are able to form the interpersonal commitments characteristic of assertion, promising, demanding, and so on.¹⁷

In order to properly understand Moran's notion of illocutionary authority, we need to first understand the notion of an illocutionary speech act.¹⁸ An illocutionary speech act, such as telling, warning, inviting, or promising, is something a speaker does when she means to perform the act in question (telling, warning or whatever), and is recognized by the audience as performing that very act with her words. This can be contrasted with her bare utterance (the ‘locutionary act’), as well the various things she may accomplish *by* or *through* speaking, such as persuading or seducing her audience (the ‘perlocutionary act’). So, if I say “it is 7 o'clock”, meaning to *tell* you that that it is 7 o'clock, and you recognise that this is what I am doing (telling you something), then I have succeeded in telling you. The same goes for inviting, warning,

¹⁷ Moran (2013, 2005)

¹⁸ Austin (1962), Searle (1969).

promising, and many other illocutionary acts: once you correctly recognise that I am inviting, warning or promising you, then that is precisely what I have done.¹⁹ Obviously, it is another thing entirely whether you believe what I tell you, accept my invitation, heed my warning, and so on, but, as Moran points out, the question of whether to believe me, accept my invitation, heed my warning (etc) does not so much as arise unless you correctly recognize what I am up to with my words.

The key point in all of this is that there is a distinctive *up-to-me-ness* or authority involved in the performance of illocutionary speech acts. It is up to the speaker herself what illocutionary act she performs, if any, and it is this authority that the audience credits her with when she recognizes her as performing the act in question. And the fact that she relies in this way on her audience's recognition reveals something about the nature of this authority. The speaker's illocutionary act is not up to her in the sense that she makes it happen, as it is with her locutionary acts, nor even in the sense that she *tries* to make it happen, as it is with her perlocutionary acts. Instead, her illocutionary speech acts are up to her in the sense that they depend on "the authority of speaker to constitute his utterance as an illocution of some kind, and hence to make it count as, say, a statement rather than a question, and thus as a possible object of belief" (Moran 2013, 123). And that is just to that power involved in illocution is an *authorial power*.²⁰

Moran cashes this idea out further by claiming that illocutionary acts are all ways in which the speaker commits herself, and thereby makes herself responsible, to her audience:

"In conferring a certain illocutionary status on his words, the speaker deliberately and explicitly alters the status of his action so as to expand or restrict the scope of his accountability to the other person" (Moran 2013, 129).

This may seem similar to Pettit's account, which also involved the speaker making herself responsible to her audience, in the sense that she makes herself liable to rebuke or sanction if she fails to maintain and manifest the attitudes she represents herself as having. But the sort of commitment and responsibility that Moran sees as involved in illocution is not of this attitude-

¹⁹ According to Jennifer Hornsby's influential account, illocutionary acts are those speech acts for which recognition from the audience is sufficient for success, i.e., they are things 'speakers can do simply by being heard as (attempting to and thus) doing them' (1994: 131).

²⁰ Tamar Schapiro (whom Moran cites) draws a helpful contrast between authorial power and empirical power: 'Whereas empirical power is the power to make things *happen*, authorial power is the power to make things *count*' (Schapiro 2001: 111).

related kind, but is instead content-related. In asserting, for example, the speaker commits herself to the truth of what she has asserted, while in promising to *phi*, she commits herself to *phi*-ing.

This last point shows how Moran's account of speaking with authority avoids the difficulties I raised for Pettit's account. On Moran's view, the authority involved in speaking for oneself is an authority to make public or interpersonal commitments, where these are understood as commitments to how things are or shall be in the world. This is exactly the sort of commitment that can underpin what I earlier called the fundamental form of discursive responsibility—the responsibility people incur in respect of the content of what they, for example, assert or promise—and so it fares a lot better than the sort of attitude-related commitment and responsibility we find in Pettit's account. Perhaps this will not appear as any great advantage to those who, like Pettit, see speech as essentially a business of speakers expressing their attitudes to others. But for those who see speech as essentially *communicative*—as our way of making claims and spreading knowledge of how things are in the world—I think Moran's approach will appear more attractive.

3.3 Groups as deliberators and illocutionary speakers

Finally, let us consider whether Moran's approach to making up one's mind and speaking with authority might be usefully extended to social groups—i.e., whether groups can author their attitudes through deliberation, and speak with illocutionary authority.

We can begin with the question of whether groups can make up their minds in the sense of deliberation culminating in judgment. In order to count as making up one's mind in this way the group must be able to address itself, in a deliberative spirit, to certain questions about the world—questions about what is true, or desirable or worth doing. And it must be recognized, both by the group itself, and by its interlocutors, that by settling these questions for itself the group immediately and directly determines its attitudes. It cannot be that the group deliberates and, say, concludes that *p* is true, yet still needs somehow to 'produce' in itself the collective belief that *p*, or rid itself of the belief that not-*p*. Rather, the collective belief that *p* must be the direct upshot of the group's deliberation that concludes that *p*.

What sort of processes or organizational arrangements might allow a group to make up its mind in this sense? I do not think the sort of sophisticated aggregation procedures favoured by Pettit, such as the straw vote, will be up to the task. This is because these processes invariably delegate

deliberation, reasoning and judgment to the individual group members, who are required to make up *their* minds on the issue at hand and then vote in line their judgment.²¹ The constellation of attitudes that constitute the ‘group mind’, then, is only an aggregative function of the views of the group members, albeit not a straightforward proposition-by-proposition aggregation function, on pain of susceptibility to the discursive dilemma. Still, it seems that even sophisticated voting systems like the straw vote procedure would not count as genuine group deliberation, in Moran’s sense of deliberation. This is because they are not a matter of the group considering and settling for itself a question of what is true, or worth doing, or desirable in the world, but rather a matter of the group members regulating the group-level ‘attitudes’ arrived at by voting.

To be clear, the reason why I think an aggregative procedure, such as a straw vote, could not amount to group deliberation is not because such a procedure is carried out by individual group members. On the contrary, I am prepared to accept Pettit’s basic claim that all group action, including group deliberative action, must involve, or be in some sense be carried out by, individual agents. The problem is not about who carries out an aggregative procedure but rather what sort of procedure it is. Deliberation involves forming or revising an attitude by considering and settling for oneself a question about the world, such as whether or not *p*, on the basis of considerations that bear on that question. But a group that arrives at its attitudes by aggregating individual votes—even if those are representative of individual judgments, and those judgments are based on individual deliberation—cannot be said to settling for itself a question about the world, on the basis of considerations bearing on that question. This is because the primary question before the voting group is the question of what its attitudes should be, and the basis for settling that question is tallying votes, rather than consideration of reasons bearing on the content of the attitudes in question. If the ‘group attitudes’ generated via such a procedure are expressive of anything it will be the judgments of participating individuals with respect to what the group’s attitudes should be, rather than those individuals’ shared sense of what reason demands with respect to the content of those attitudes.

In contrast to the sophisticated aggregative procedures envisaged by Pettit, I suggest that what it takes for a group to genuinely make up its mind is the adoption or establishment of a plural first-person perspective, or shared rational point of view.²² A set of people are able to participate in the establishment and maintenance of such a perspective, but this is not a matter of their finding

²¹ See Rovane (2014) for a similar criticism.

²² Cf. Rovane (1998, 2014), and Schmid (2014, 2016).

some ‘rationally optimal’ way of aggregating what each of them already thinks, nor is it a matter of negotiating with one another over a set of attitudes each of them is prepared to accept and live up to ‘in the name of the group’. Such posturing is not the equivalent, at the collective level, of genuinely making up one’s mind, but is at best a matter of emulating or masquerading as a person. This is because it does not involve their self-consciously committing themselves on how things are in world—on what, as far as they’re concerned, is true, or good, or worthwhile in it. When we make up our collective mind we cannot take ourselves to be constructing, attitude by attitude, an ‘outlook’; rather we must *look out* together and form our commitments based on what we see.

The kind of co-deliberation that is needed is not restricted to formal organisations and institutions, but can be achieved even by small, unstructured groups. Consider, for example, my wife and I, who face the question of how much screen time our child should be allowed per week. This is a question that arises for *us*, since it relates to something we do together, namely, parenting our child. It may well be that were I to consider and settle that question for myself I would arrive at a different answer from the answer my wife arrives at, so that we would need to try to change each other’s mind, or find some kind of compromise. But of course we need not approach the question in that way, as separate deliberators. Instead, we could make the deliberation itself a joint activity, figuring out together what evidence bears on our question, what epistemic standards we should employ, and then, using those standards, what conclusion that evidence supports. If we do this, and, say, conclude that no more than two hours screen time per week is acceptable, then we have therein jointly authored a collective attitude, the belief that no more than two hours screen time per week is acceptable. In so doing we have not acted on ourselves in order to somehow ‘produce’ a collective belief; rather, we have invoked an authorial power, the power to make this judgment, that no more than two hours screen time per week is acceptable, count for us as our belief. This kind of making up our mind is not some sort of higher grade reflective capacity that can be emulated only by suitably designed organisations; it is a ubiquitous and everyday feature of shared human life.²³

²³ To a certain way of thinking about collective phenomena, this suggestion may appear thoroughly unilluminating, perhaps even circular. After all, the question is how a group can deliberate, and the answer I’ve proposed, in short, is that the group must do what deliberation requires: *we* must *together* consider and settle certain questions for *us*. But shouldn’t I be dispensing with such ‘we’-talk if I am to offer a genuine explanation of what happens when a group does this—what is actually going on in such cases? I do not think so. Indeed, I think it should be part of the very idea of a shared rational point of view or a joint perspective that such we-talk is indispensable, or, better, that, like I-talk, it captures a basic, irreducible form of human subjectivity—plural subjectivity.

Now, of course, to eschew the demand for an ultimately individualistic explanation is not to deny that a lot more does need to be said about just what such plural subjectivity consists in, and how it compares to individual subjectivity. I cannot hope to provide a full characterisation here, but one way of fleshing it out, that I think would

Similarly, when speaking for oneself is understood as the exercise of illocutionary authority, there is little mystery as to how it could be accomplished by a group. Or perhaps it would be better to say there is no *further* mystery as to how it could be done by groups, because it involves the very same sort of authorial power that is involved in the individual case. This is not the power to self-represent with confidence, as Pettit's account has it, but rather the power to make one's words count in certain ways, and therein make commitments and undertake responsibilities to others in respect of how things are in the world. This is something that groups manifestly do.

To see this, note that there are many different ways in which a group can 'speak' in the sense of coming out with meaningful utterances. These may include a spokesperson's declarations, but they also extend, for instance, to co-authored documents, online petitions, the chants of protestors, and even collective gestures of various kinds. Regardless of their particular locutionary instruments, it is clear that when groups speak they also typically mean to have their locutions count in certain ways—count, for example, as asserting, or as promising, or as demanding. And not only do groups of various shapes and sizes lay claim to this kind of authorial power, they are also typically credited with it. As readers, we understand that the co-authors are *asserting* something; as members of the public, we recognize that the government is

complement the broadly Moran-ian approach advocated here, has been developed in recent work by Hans Bernhard Schmid (Schmid 2014, 2016). Schmid's basic idea is that plural subjectivity, like individual subjectivity, centrally involves a distinctive kind of self-consciousness or self-awareness. Just as an individual's conscious attitudes are immediately known to her *as hers*—this being an essential aspect of having a first person perspective at all—so too there is a way of sharing attitudes with others such that they are known by those sharing them *as theirs*. If you and I discuss the question of whether good bagels can be found in Vienna, and we conclude that they cannot, then the shared belief we arrive at, that good bagels cannot be found in Vienna, is one we will be aware of *as ours*, as opposed to just a belief we happen each to have arrived in the course of our conversation. And, similarly, just as having a first-person perspective on one's attitudes involves knowing that one is subject to the normative pressures associated with being the bearer of those attitudes—the pressure to ensure their overall coherence, the pressure to infer what they jointly imply, and the pressure to justify them if challenged—so too it appears there is often a shared awareness of various normative pressures associated with the establishment of shared attitudes. If you and I are discussing whether good bagels can be found in Vienna, and you think they can while I think they cannot, then each of us will feel a kind of normative pressure to resolve our disagreement and come to a common view. I cannot pursue it any further here, but it does seem plausible that the sort of group deliberation I have tried to highlight, in contrast to the aggregative procedures of group decision-making favoured by Pettit, would be rooted in some kind of plural self-awareness.

Still, it should be acknowledged that even if appeal to a notion of plural self-awareness goes some way towards characterising the indispensable 'we' involved in group deliberation, many thorny questions remain. Does a lack of agreement amongst individuals preclude the formation of a group attitude? If so, then it appears that group attitudes will be very hard to come by; but if not, then it is unclear just how much agreement is needed, and also how groups can avoid the sort of profiles of individual views that give rise to the discursive dilemma. Moreover, we might wonder what sort of psychological processes are involved in deliberating from a shared point of view? If these are conceived individualistically, as the capacity of individuals to adopt a 'group perspective', then it will of course be possible for some group members to do so while others do not. Is there genuine group deliberation in such a case? These are important questions, but I will have to set them aside in this paper. I thank an anonymous reviewer for bringing these and other difficulties to my attention.

promising us, through its spokesperson, that stricter gun control measures will be introduced; and as onlookers, we see that the group of protestors is *demanding* fair and equal treatment.

It is in these ways that groups are able to forge the sort of discursive commitments that are properly seen as distinctive of personhood—commitments made to an audience, but in respect of some feature or fact in the world. When the co-authors assert that *p* they therein commit themselves to the truth of *p*, making themselves answerable to their audience for challenges to the truth of *p*. When the government promises stricter gun control they therein commit to providing it, making themselves liable to rebuke or sanction from the public should they fail to do so. And when the protestors demand fair and equal treatment, they make a claim to the effect that this is something owed to them, thus making themselves answerable to others for this claim. Like the ability people have to co-deliberate and so jointly make up their mind, this ability to forge collective commitments through speech is also a ubiquitous feature of shared human life.

Conclusion

My primary aim in this paper has been critical. I have tried to challenge Pettit's claim that groups qualify as persons because they can "make up their minds" and "speak for themselves". I do not deny that some capacity for "making up one's mind" and "speaking for oneself" is central to qualifying as a person, nor do I think that Pettit has failed to show that groups can do what *he* thinks these capacities involve. Instead, my complaint is with how he conceives of these two distinctively personal capacities. Specifically, I argued that the way he construes making up one's mind fails to do justice to the first person perspective, and the way he construes speaking for oneself fails to do justice to our interpersonal discursive practices of holding one another responsible in respect of the contents of what we say. So Pettit fails, I think, to make a solid case for collective personhood because his model of personhood is itself inadequate.

In the final part of the paper I tried to supplement this negative conclusion with something more positive: a brief sketch of Moran's alternative approach to "making up one's mind" and "speaking for oneself" and a suggestion that it might be extended to groups. Of course, a great deal more needs to be said both in defence of this alternative proposal and by way of its application to situations in which people deliberate together and speak with one voice. I shall leave the development of the proposal for future work; here, I hope simply to have shown it to be a proposal worth developing.

Acknowledgments

For helpful comments on earlier versions of this paper, I am grateful to Olav Gjelsvik, Jennifer Hornsby, Sarah Stroud, Philip Pettit, Hans Bernhard Schmid, and Dina Townsend, as well as two anonymous reviewers for the *Journal of Social Philosophy*. My research is funded by the Austrian Research Council (FWF), grant number I 3068.

Works Cited:

Austin, John. L. 1962. *How to Do Things with Words*. Oxford: Clarendon Press.

Bilgrami, Akeel. 2006. *Self-knowledge and resentment*. Harvard: Harvard University Press.

Boyle, Matthew. 2011. "Making up Your Mind' and the Activity of Reason". *Philosophers' Imprint* 11 (17): 1-24.

Gilbert, Margaret. 1989. *On Social Facts*. London: Routledge.

- 2014. *Joint Commitment: How We Make the Social World*. New York: Oxford University Press USA.

Hieronimi, Pamela. 2006. "Controlling attitudes." *Pacific Philosophical Quarterly*. 87 (1): 45-74.

- (2009). "Two kinds of agency." In O'Brien, Lucy and Soteriou, Matthew (eds.), *Mental Action*. Oxford: Oxford University Press: 138–162.

Hornsby, Jennifer. 1994. "Illocution and its significance." in Tsohatzidis, Savas (ed), *Foundations of Speech-act theory: Philosophical and Linguistic Perspectives*. London: Routledge.

List, Christian. 2014. "Three Kinds of Collective Attitudes." *Erkenntnis* 79 (9):1601-1622.

- 2016. "What is it Like to be a Group Agent?" *Notis* 52:2: 295-319.

List, Christian and Pettit, Philip. 2002. "Aggregating sets of judgments: An impossibility result." *Economics and Philosophy* 18 (1):89-110.

- 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.

McDowell, John. 1998. *Meaning, Knowledge, and Reality*. Harvard: Harvard University Press.

McGeer, Victoria. 1996. "Is "self-knowledge" an empirical problem? Renegotiating the space of philosophical explanation." *Journal of Philosophy* 93 (10):483-515

- 2008. "The moral development of first-person authority." *European Journal of Philosophy* 16 (1): 81–108.

McGeer, Victoria and Pettit, Philip. 2002. "The Self-regulating Mind." *Language and Communication* 22: 281-99.

Moran, Richard. 1999. "The authority of self-consciousness." *Philosophical Topics* 26: (1-2): 179-200.

- 2001. *Authority and Estrangement: An Essay on Self-Knowledge*. New Jersey: Princeton University Press.
- 2005. "Getting told and being believed." *Philosophers' Imprint* 5:1-29.
- 2012. "Self-Knowledge, 'Transparency', and the Forms of Activity". In Smithies, Declan & Stoljar, Daniel (eds.), *Introspection and Consciousness*. Oxford: Oxford University Press.
- 2013. "Testimony, Illocution and the Second Person." *Aristotelian Society Supplementary Volume* 87 (1):115-135.

Pettit, Philip. 2001. *A Theory of Freedom: From the Psychology to the Politics of Agency*. Oxford: Oxford University Press.

- 2003. "Groups with minds of their own." in Schmitt, Frederick. (ed.) *Socialising Metaphysics: the nature of social reality*. London: Rowman & Littlefield Inc.
- 2007. "Rationality, Reasoning and Group Agency." *Dialectica* 61 (4):495-519.
- 2009. *Made with Words: Hobbes on Language, Mind, and Politics*. New Jersey: Princeton University Press.
- 2012. *On the People's Terms: A Republican Theory and Model of Democracy*. Cambridge: Cambridge University Press.

- 2014a. "How to tell if a group is an agent." in Lackey, Jennifer (ed.) *Essays in Collective Epistemology*. Oxford: Oxford University Press.
- 2014b. "Group Agents are Not Expressive, Pragmatic or Theoretical Fictions." *Erkenntnis* 79 (9):1641-1662.
- 2016a. "Making up your mind: how language enable self-knowledge, self-knowability and personhood." *European Journal of Philosophy* 24(1):3-26.
- 2016b. "The Birth of Ethics" in Matheson (ed), *The Tanner Lectures on Human Values*, Vol 34. Salt Lake City: University of Utah Press.

Pettit, Philip & Schweikard, David. 2006. "Joint actions and group agents." *Philosophy of the Social Sciences* 36 (1):18-39.

Popper, Karl. 1957. *The Poverty of Historicism*. London: Routledge.

Rovane, Carol. 1998. *The Bounds of Agency: An Essay in Revisionary Metaphysics*. New Jersey: Princeton University Press.

- 2014. "Group Agency and Individualism." *Erkenntnis* 79 (9):1663-1684.

Schapiro, Tamar. 2001. "Three conceptions of action in moral theory." *Noûs* 35 (1):93–117.

Schmid, Hans Bernhard. 2014. "Plural self-awareness." *Phenomenology and the cognitive sciences* 13 (1):7-24

- 2016. "On Knowing What We Are Doing Together." in Brady & Fricker (eds.) *The Epistemic Life of Groups: Essays in the Epistemology of Collectives*. Oxford: Oxford University Press.

Searle, John. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge: Cambridge University Press.

Tollefsen, Deborah. 2002. "Organizations as true believers." *Journal of Social Philosophy* 33 (3):395–410.

- 2015. *Groups as Agents*. Cambridge: Polity.

Tuomela, Raimo. 2013. *Social Ontology: Collective Intentionality and Group Agents*. Oxford: Oxford University Press.

Watkins, John. 1957. "Historical explanation in the social sciences." *British Journal for the Philosophy of Science* 8 (30):104-117.