



universität
wien

MASTERARBEIT

Titel der Masterarbeit:

Can We Know?

Justified Group Belief and Social Ontology

Verfasserin:

Katharina Maria Bernhard, BA

Angestrebter akademischer Grad:

Master of Arts (MA)

Wien, 2015

Studienkennzahl:	066 941
Gebiet lt. Studienblatt:	Masterstudium Philosophie
Betreuer:	Univ. Prof. Dr. Hans Bernhard Schmid

Acknowledgements

There still exists the cliché of the philosopher sitting alone in an armchair, thinking deeply and waiting for an original idea (or argument) to strike her mind. But things do not usually work that way. While there is indeed a lot of sitting in chairs, reading and thinking deeply involved, bigger projects like a Master's thesis are rarely a completely individual endeavour. After all, philosophers are, like all humans, social beings. Hence there exists a wide range of people to whom I am indebted for their support in the production of this thesis, not only directly by supervising it and/or giving valuable feedback, but also in many indirect yet equally essential ways. Moreover, there exist two groups that provided a most vibrant climate for the completion of such a project, thereby influencing this thesis in both direct and indirect ways.

First, I wish to thank every participant of the social ontology workshop at the Department of Philosophy at the University of Vienna. Among those are Bernhard Schmid, Alessandro Salice, Thomas Szanto, Gerhard Thonhauser, Michael Schmitz, Odin Kröger, Gerhard Kreuch and Gloria Mähringer. I benefitted greatly from discussing a twenty-page draft of this thesis in the course of one of our meetings. Most importantly, however, my deepest gratitude goes to the initiator of this group, Hans Bernhard Schmid, who also supervised this thesis.

In fact, mere thankfulness cannot cover what I owe to Bernhard Schmid. In addition to his supervision of this thesis, I have most benefitted from his vast expertise in social matters, both philosophical and practical over the last few years. Working for Bernhard Schmid has been an invaluable experience and if this thesis is of any quality, it is due to him. I still do not know what makes me deserving of his trust and kind support – with respect to this thesis as well as with regard to my studies in general – but I am most honoured by and deeply grateful for it.

Furthermore, I am grateful to Thomas Szanto who read through parts of a draft version of this thesis. His feedback and criticism were most valuable. Also, I had the opportunity to discuss various ideas for some of this thesis' discussions with him and I am particularly grateful for his valuable comments and willingness to listen.

Moreover, I am indebted to Gloria Mähringer. We were in the same boat when writing this Master thesis and accordingly, we could relate to one another's struggles. I consider myself lucky to call Gloria a friend.

The second group that I wish to thank is the Vienna Forum for Analytic Philosophy (WFAP). Dejan Markovec, Günther Eder, Karoline Paier, Lukas Heuberger, Katharina Sodoma, and many others read through the above-mentioned twenty-page draft, which I had the opportunity to present twice, in two different but equally confused versions, in our Friday sessions. Most importantly, however, there exist two members of the WFAP to whom I am particularly grateful.

Sebastian Kletzl, chair of the WFAP and dear friend, read through a draft version of this thesis thoroughly and provided extremely helpful comments and criticism. Moreover, he was one of the very few people I would regularly discuss my thesis with. It was not only Sebastian's expertise regarding social epistemology, but also his ability to somehow realize, even when I did not, that I was in need of some uplifting words that proved essential in the production of this thesis. This thesis clearly carries the markings of our discussions.

In addition, this thesis owes a lot to Patrick Johannes Klug's successful struggle to teach us, members of the WFAP, the basics of Bayesian Epistemology in the course of nearly two semesters. Patrick's competent, thoroughly designed, and always-amusing lectures were an enormous help in enabling me to provide clarity in the more technical side of my discussions. He too read through parts of a draft version and his criticisms and comments were most valuable. I greatly appreciate his efforts.

It is not only professional relationships or instances where they coincide with friendships that shape such a project or significantly contribute to its successful production. In our joint lives, my husband Christoph and I try to cover each other's backs and through the entire time of writing this thesis, this relationship was clearly one sided and to my advantage. In particular, he always put my struggles first, despite having enough to deal with on his own. I hope to be able to balance this disparity some day.

Furthermore, I owe my deepest gratitude to Christine, my mother, who relieved me of various obligations that I would have had to face over the last few years, thereby enabling

me to focus on my studies. She also covered my back (and still does) and I am most thankful for this.

Finally, writing in a non-native language is always a challenge. I am thankful to Caroline Amy Peña Bray for her excellent work in proofreading this thesis, and in particular for her patience in correcting the same grammar and spelling mistakes again and again.

Contents

Introduction	4
PART I	
Chapter 1 Two Preliminary Assumptions	16
1.1. Non-Summativism about Groups	17
1.2. Epistemic Externalism and Process Reliabilism	24
Chapter 2 Towards an Epistemology of Groups. Some Essentials	32
PART II	
Chapter 3 List and Pettit's Group Agents	40
3.1. Agency and Rationality: The Positive Claim	43
3.2. A Theory of Attitude Aggregation: The Negative Claim	46
Chapter 4 Meeting the Knowledge Challenge. List and Pettit's Proposal	57
4.1. Truth-Tracking and Probabilities of Judgement	60
4.2. Epistemically Promising Strategies of Judgment Aggregation in Groups	65
Chapter 5 Externalism, Group Agents, and Two Forms of Reliabilism	71
5.1. Epistemic Externalism and Group Agents. Where the Problem Does Not Lie	74
5.2. Truth Tracking. A Potential Problem	82
5.3. Reliability of Belief vs. Reliability of Judgements. A Problem	86
PART III	
Chapter 6 Adopting a Process Reliabilist Framework	94
6.1. Goldman's Social Process Reliabilism	96
6.2. Automatic Transfer of Justification. Two Models	101
Chapter 7 One Problem for Social Process Reliabilism and Two Solutions	105
7.1. Saving List and Pettit's Group Agents. A Non-Reductionist Inter-Personal View	113
7.2. Saving the Intra-Personal Model. Rovanean Group Persons	118
Summary and Conclusion	132
References	140
Appendix	146

Introduction

This thesis is concerned with collective epistemology and, in particular, group beliefs and knowledge. In her “Collective Scientific Knowledge” (2012) Melinda Fagan states that there exist multiple ways in which knowledge can be collective in nature. We can assume that (a) a belief is collective if its production necessarily involves a group or collective effort, or a shared framework of some sort. However, this is not the same as saying that (b) a belief is held by a group itself (as opposed to it being held by a majority of its members or a relevant representative of the group and so on). While these two ways in which beliefs can be collective do not exclude each other, the approach required for (a) is quite distinct from that required for (b). Moreover, (a) does not necessarily entail (b). Indeed, there exists a vast amount of literature that addresses the various ways in which knowledge, and in particular scientific knowledge, is necessarily collective with regard to (a).¹ Alternatively, this thesis is concerned with collective belief in terms of (b), and understands (a) to come from claiming (b). More precisely, it is concerned with the question of how a group, understood as a single entity independent of its members, can hold justified beliefs or even knowledge.

Asking how a group can hold a justified belief touches upon two philosophical areas that are often kept apart in most discussions, namely social ontology and (social) epistemology. Social ontology is essential to understanding what a group, taken as a single independent entity, is. This is necessary when working to comprehend the idiosyncrasies of collective attitudes when compared with individual attitudes. In this study, the job of epistemology is to give an account of whether or how the beliefs held by a group can be epistemically justified or constitute group knowledge in context of the social ontological underpinning.

Although nearly all ontologists and epistemologists agree that human beings exist and that they have certain mental states, like beliefs, desires or intentions, such a comparable

¹ There exists a wide discussion regarding the social dimension of scientific knowledge that takes on various forms and deals with various issues. Moreover, it covers a wide range of different disciplines in addition to philosophy and the philosophy of science, such as, sociology, history, economics, anthropology and ethnology. For a brief historical overview regarding the emergence of debating knowledge as a social phenomenon in the western hemisphere, see Burke (2000; 2013). For a brief overview of more contemporary debates in a rather systematic manner, see Longino (2013).

level of agreement remains absent from both social ontology and social epistemology when dealing with group phenomena. In fact, whether groups exist and, if they do, what their status is (both ontologically and epistemologically) are controversial matters. Thus the ontological grounds of an endeavour in group epistemology need be agreed upon before entering into a discussion. In light of this, one particular challenge stems from the fact that standard epistemology is usually only concerned with individual subjects, hence the tools we make use of are made to measure for human beings. Taking into consideration the discrepancy between epistemology's rather individualistic tool box and the requirements imposed on an epistemological discussion given the specifics of group belief, a central question appears: how could or should we adjust these tools? This thesis discusses and critically examines a number of suggestions concerning how to best adjust our epistemological framework.

To the degree that standard epistemology distances itself from discussions of collectives and knowledge familiar to the philosophy of science and adjunct areas, the task of dealing with group phenomena, and in particular with (b), is a rather new phenomenon in "standard" epistemological terms.² Accordingly, there exists a considerable amount of uncharted territory regarding the above-mentioned questions, and this thesis takes steps to acknowledge and account for this fact. Although the literature on this specific matter is limited, a general overview of the problem is not within the remit of this thesis. However, it will be important to hone the focus of the present study. With regard to groups and group knowledge, delineating what these two terms amount to is key. In agreement with recent debates, this thesis presupposes two general claims, one with respect to social ontology and the other with respect to epistemology. While I introduce these choices, I do not argue for these claims in their own right, that is to say, I do not defend them in greater detail against rival claims within social ontology or epistemology, nor do I discuss general objections (unless a moment arises at which it is necessary to define their essential features). It is worth noting that the

² Examples of the current state of the collective epistemology debate along rather "standard" epistemological lines encompass, for example, Schmid, Sirtes, and Weber (2011); Mathiesen (2006, 2007); Bird (2010, 2014); Tuomela (2004); Hakli and Negri (2011); Hakli (2007); Thagard (1997); Corlett (2007); Solomon (2006); Gilbert (1987, 2004); Goldman (1999, 2004); Kopec (2012); Lahroodi (2007); Lackey (2014); List (2005); Schmitt (1994); Wray (2001, 2010); Tollefsen (2004, 2006). In particular, discussions regarding groups have entered into recent debates via the topic of testimony from groups, see in particular Lackey (2014a); M. Fricker (2012); Tollefsen (2007).

discussions within this thesis might be best understood as being based on two conditional clauses. First, if the social ontological claim were true and second, if the epistemological framework were correct, how would things stand with respect to justified group beliefs or even knowledge?

The ontological claim relates to the types of group that this thesis is concerned with and amounts to what could be called “non-summativism” about groups.³ In brief, the idea of groups discussed here is that groups constitute an entity in their own right, where group attitudes are not reducible to those of its members. The primary focus herein is that of a particular representation of non-summativist accounts, namely Christian List and Philip Pettit’s group agents. The epistemological claim is concerned with the epistemological framework under which the possibility of group knowledge (or justified group belief), given the ontological claim, is discussed in this study. Here the thesis is committed to externalist accounts and, in particular, to process reliabilism. These two claims are established in the first chapter of this thesis. With the necessary restrictions to hand, the paper then discusses some general implications for group epistemology and in doing so highlights what may indeed constitute the area’s most pressing issue (Chapter 2).

I do not discuss each of these premises in their own right for several reasons. First, such a project would exceed the aim of this thesis, or rather, make it impossible to reach the aim of discussing the possibility of justified group beliefs altogether. More importantly, and at least with respect to the chosen epistemological framework, it would prove particularly difficult for me to provide such a defence due to the very motivation behind this thesis. As motivations have it, they are pragmatic rather than philosophical. Yet I do hope that they contain a certain amount of philosophical value.

A vast amount of the literature covers the topic of collectives and knowledge, especially in the philosophy of science and adjunct areas. Focus usually rests upon an explanation of how justification is a group-internal issue, depending on paradigms or styles of thought;⁴ on the extent to which being a member of a certain (scientific) group allows the said members to know something; or on how knowledge in a group (as

³ See Chapter 1 for further discussion.

⁴ See, in particular Kuhn (2012), Fleck (1980).

opposed to “of” a group) might emerge given various forms of division of labour. Hence they are mainly, though not exclusively, concerned with (a), that is to say, with the ways in which scientific knowledge is necessarily collective or dependent on collective aspects. These are pertinent questions and the project of this thesis falls close to the aforementioned discussions, at least to the extent that (b) is accompanied by (a). Yet, although there exist a lot of interesting discussions in these debates that could certainly be made fruitful for engaging in epistemological discussions regarding group phenomena, this thesis’ focus is slightly different to these discussions. Epistemologists such as Alvin Goldman, Jennifer Lackey, or Duncan Prichard (representatives of the “standard” reliabilist epistemology branch) do not, or only to a very limited extent, engage with these areas.

However, at the same time there is an on-going attempt, most prominently defended by Alvin Goldman, to socialise epistemology at least to a certain degree (see Goldman 2010). Goldman seeks to defend a view in which “collective doxastic agents” (ibid, esp. pp. 25ff) are a beneficial object of epistemic inquiry. He even expands epistemology’s concern to topics such as the acquisition of social evidence or the evaluation of the epistemic success of social institutions such as courts. Such projects qualify as standard epistemology to the extent that they remain in tune with certain features of (analytic) epistemology. These features include:

(A) The epistemic agents of traditional epistemology are exclusively individuals. (B) Epistemology focuses on the study of epistemic evaluation or normativity, represented by such evaluative concepts as justifiedness, rationality, and knowledge. Traditional epistemology asks such questions as how individuals can acquire knowledge and maintain justified or rational credal states. (C) Traditional epistemology assumes that the normative standards of rationality and justifiedness are not merely conventional or relativistic, but have some sort of objective validity. (D) The central notions of epistemic attainment – knowledge and justification, for example – either entail truth or are closely linked to truth. A known proposition must be true; and justified beliefs, according to many mainstream views, are in some sense likely to be true. (E) Truth is assumed to be an objective, largely mind-independent, affair. (F) Traditional epistemology takes its central business to be the critical examination of doxastic ‘decision-making’ (DDM): adopting, retaining, or revising one’s beliefs and other doxastic attitudes. (Goldman 2010, p. 2)

While Goldman claims to remain within these cornerstones of epistemology, according to these criteria there exists a wide range of theories that do not count as “epistemology”. In particular, Goldman argues that what he calls “revisionist” theories should be rejected

as “real” epistemology. Revisionist theories involve “postmodernism, deconstructivism, social constructivism, and various social studies of science, including the ‘strong programme’ in the sociology of science” (ibid p. 3).⁵ Since such approaches breach most or perhaps all of the criteria posed in (A)-(F), they are disqualified from being epistemology or even social epistemology. The other two approaches to social epistemology, of which Goldman conceives himself to be a part, are preservationism and expansionism. The first seeks to adhere to all of the standard criteria for (analytic) epistemology, while allowing for certain social phenomena to be proper objects of investigation within such a framework. Expansionism seeks to expand to disciplines other than epistemology while not violating these norms either (but probably challenging them). In this sense, Goldman contests, his process reliabilist approach to social epistemology can be labelled “real” epistemology. This is so even though his assumption that collective doxastic agents constitute a proper object of epistemological inquiry appears to be in some tension with (A) (see Palermos and Pritchard 2013). Yet despite this tension (which may be less dramatic if we depart from the idea that “individuals” as mentioned in (A) need to be human individuals), Goldman remains committed to individualism in, more or less, all other respects and in particular remains sceptical about other ways in which the social may play an essential role in knowledge and knowledge production.⁶ At its core, therefore, Goldman’s process reliabilism remains individualistic, to the extent that knowledge is something an individual (human being or group) needs to arrive at by means of processes that take place only in the relationship between the world and an individual’s cognitive processes.⁷ This commitment has been challenged from within “standard” epistemology by, for example, Orestis Palermos and Duncan Pritchard’s concept of extended knowledge (Palermos and Pritchard 2013). Another

⁵ Some theories Goldman probably has in mind here are Foucault (1972); Bourdieu (2004); and various post-structuralist theories as well as, sociologies of knowledge such as Bloor’s (1991, 2005), for a more recent discussion in line with this approach, see Kusch (2002). Also, there exist various contributions in context of feminist epistemology which probably too disqualify as epistemology according to this proposal (see Alcoff and Potter 1993 for an example).

⁶ For example, Goldman is a strict reductionist regarding testimony. Furthermore, he defends (C) and (E), which are individualist criteria, to the extent that they exclude social factors to play into both the concepts of justification and truth.

⁷ The degree to which this is a problematic assumption, at least if the other criteria are to be saved, is discussed in Chapters 6 and 7.

approach, with which this thesis is concerned in its final part, is Sanford Goldberg's anti-Process Individualism (see Goldberg 2010, 2007).

It is important to note that this thesis presupposes a "standard" epistemology, hence preservationism or expansionism, for the most part and is, as thus, committed to the above criteria at the expense of considering any broader social context that may play into the production of knowledge. It only departs from (A) to the extent needed for it to allow for collective agency and leaves, for example, (C) and (E) untouched.⁸ Yet this does not mean that this choice reflects the personal preferences or philosophical commitments of its author, as Chapter 7 will indicate, where discussion indeed departs from some of these criteria. On the contrary, my choosing of a standard epistemological framework, and especially that of process reliabilism, is motivated by a certain unease with respect to the idea that, to adopt Goldman's vocabulary, a social epistemology that allows for group agents can circumvent even modest forms of revisionism, especially with respect to feature (A). Frankly, I doubt that non-summativism about groups – even in rather modest versions, such as List and Pettit's group agents – can be combined successfully with an epistemological framework in which, at the end of the day, things boil down to individuals and their cognitive processes or in which collective issues can form debates analogous to (human) individual epistemologies. Accordingly, the idea behind this thesis is to provide a detailed and critical discussion of this line of thought in order to see whether it really does provide us with a workable epistemological framework for making sense of justified group belief or perhaps even group knowledge. I doubt that this project would be possible without abandoning more individualist commitments than the authors discussed are probably willing to let go of. The last two chapters of this thesis will provide reasons in support of this doubt. Accordingly, the restrictions that this thesis forces upon itself stem from a certain sympathy for the very camp that I neglect to discuss and which I will, for most part, completely omit.

One consequence of the tensions between the chosen framework and my actual philosophical commitments is that I often refrain from making objections where I actually believe an objection is due. For example, there is a lot that could be said about

⁸ For more recent discussions regarding how to challenge other criteria in the above-mentioned list within epistemology, see for example Wright (2006); García-Carpintero and Kölbel (2008); Kölbel (2004); Hales (2011).

the idea of “collectivity” and the joint attitudes that List and Pettit stress. Instead of criticising their account right away, however, I choose to buy into most of their claims in order to see where exactly it is that they become implausible in the course of an epistemological discussion about group belief. (As shall be demonstrated in the last chapter of this thesis, Goldman’s social process reliabilist picture of epistemic justification of group beliefs goes against its own individualistic requirements.)

Restrictions on the epistemological side are equally difficult to come by.⁹ In particular, there are various reasons to question not only (A), but also (C) and (E). Furthermore, there exist numerous well-known objections to or worries about the highly idealized frameworks this thesis will discuss, in particular concerning Bayesian epistemology and judgement aggregation theory.¹⁰ Both of these approaches are particularly important for List and Pettit’s account of group agents, as well as Goldman’s social process reliabilism. Possible objections target various issues, from the very concept of rationality that is presupposed in both fields, to the problem of the priors. The latter problem, for example, makes it difficult, at least for some philosophers, to accept Bayesian epistemology as a proper or even possible framework for an analysis of justified belief formation. Disagreeing with certain assumptions regarding epistemic justification or the plausibility of truth-tracking accounts may also lead to the dismissing of a theory. But it may just as well lead to being rejected by that theory, which is one reason why I refrain from making such general objections and remain fairly neutral regarding the general plausibility of such approaches.

Instead of doubting each of the theories discussed herein in their own right, this thesis aims to show that even if we accept the claims these theories make or presuppose, their aim to provide an epistemology of groups that preserves certain aspects of (A), while at the same time allowing for non-summativism about groups, is unlikely to succeed. To make such a claim, however, this thesis will have to side with both, the social ontological claim as well as the epistemological one for the most part. Yet, this does not mean that by restricting the focus in that manner, no objections could be made. Rather, it only restricts

⁹ Even more so, since this thesis’ author considers herself a revisionist in Goldmanian terms.

¹⁰ See, for example, Hájek (2008) and (Rovane 2014).

the kind of objections that are sensible in this context to those that regard claims concerning the combination of the ontological side and the epistemological one.

It remains an open question, however, as to whether this project succeeds or is able to succeed in providing objections that, at the very least, follow the rules of the particular debates and theories discussed in this thesis. After all, what Laurence Bonjour diagnoses with respect to epistemic externalism might be just as apt an observation for the presupposed concepts of rationality, Bayesian ideal agents or judgement aggregation theory:

The problem [...] is that this very radicalism [of epistemic externalism] has the effect of insulating the externalist from any very direct refutation: any attempt at such a refutation is almost certain to appeal to premises that a thoroughgoing externalist would not accept. (Bonjour 2001, p. 14)

In light of Bonjour's observation, this project may be optimistic at best and utterly naïve at worst. I prefer to opt for the former, and hence prefer to believe that it succeeds in showing that at least some of the elements of ontological and epistemic individualism that the accounts discuss presuppose, and moreover attempt to save in their discussion of groups and group beliefs, are not sustainable.

This project proceeds as follows. In the first part (Chapters 1 and 2), I outline the working hypothesis of the thesis. In Chapter 1, the two preliminary assumptions regarding both the social ontological and epistemological underpinnings are introduced. One side of social ontology, namely non-summativism about groups, is assumed and briefly introduced. As we shall see, the term "non-summativism" used in this thesis covers more than the idea that groups exist independent of their members. It also includes non-reductionism about group attitudes as well as an assumption regarding the kinds of member involvement. On the side of epistemology, epistemic externalism is assumed. In particular, this thesis is committed to process reliabilism. Both of these assumptions should merely serve as guidance when approaching the matter of group knowledge and thus will not be presented in as fine grained a manner as a discussion of each of these claims in their own right would require.

In Chapter 2, I introduce some key issues regarding group epistemology given the two preliminary assumptions presented in Chapter 1. The relationship between member beliefs on the one hand and group belief on the other, and the emergence of a group belief given member beliefs is, in particular, identified as key to group epistemology (see Mathiesen 2006)

With this necessary background to hand, part 2 (Chapters 3 to 4) will introduce and discuss one particular account of groups, namely List and Pettit's group agents (see List and Pettit 2011). This part will be fairly descriptive in nature and should provide the necessary grounds for entering into discussions regarding the epistemic status of group beliefs. Chapter 3 introduces List and Pettit's account of group agents in general, focusing on two essential claims for the existence of group agents. These are that, first, groups can hold attitudes and perform rationally and that, second, in the course of a group's attitude formation, a group's attitudes can come apart from member attitudes in many, partly unexpected, ways. Both of these claims support the conclusion that groups exist in their own right.

Chapter 4 concerns itself with the epistemic dimension of a group's performance as an agent. According to List and Pettit (and List in individual publications), a group can, in principle, successfully track truths given the proper circumstances. This part is fairly technical, since this claim stems from certain possibility results in judgement aggregation theory. Hence some of the basics of Bayesian probabilistic theory and judgement aggregation theory (such as Condorcet's Jury Theorem) are introduced and discussed. All the same, I keep these technical discussions to a minimum and only focus on the parts that carry particularly important implications for furthering the discussion.

Chapter 5 constitutes a transitional chapter in which the epistemic implications of List and Pettit's account are evaluated and certain problems with this account are identified. In this chapter, I point to three potential problems that concern two assumptions essential to List and Pettit's epistemological framework. Since List and Pettit's proposal is also externalistic, I argue that, and explain why, these problems are specific to List and Pettit's particular version of reliabilism (which builds on a probabilistic version of Nozick's definition of knowledge as well as on judgement aggregation theory), and do not necessarily challenge a general commitment to externalism. Accordingly, this allows me

to turn to Goldman's social process reliabilism in Chapter 6, both because it is possible and because there are potential reasons for a rejection of List and Pettit's reliabilism.

Goldman's social process reliabilism indeed encompasses many valuable considerations regarding a process reliabilist framework for dealing with group phenomena. Nevertheless, is not without its problems. Chapter 7 identifies these problems and provides two routes for how to overcome them. Both of these routes enter into uncharted territory in that they combine social ontological claims with an epistemological framework in a way that has not yet been provided. Hence they should only be considered an outlook rather than a genuine contribution to group epistemology. Both of these suggested approaches avoid the problems of Goldman's social process reliabilism and involve more departures from individualistic commitments, such as in (A), than Goldman is probably willing to accept.

All in all, the discussion provided in this thesis should give reason for us to attest that indeed, something is at odds with recent discussions about group knowledge in process reliabilism, unless further concessions towards a more social picture of belief formation in groups are made.

PART I

Chapter 1

Two Preliminary Assumptions

Ontology is to epistemology what epistemology is to ontology: an unwelcome house guest, thrown out of the front door only to reenter via the backdoor. This thesis constitutes an attempt to face this situation head on and give credit to both the epistemological and ontological dimensions in as balanced a manner as possible.¹¹ This is so because I believe that both of these philosophical disciplines depend on one another, at least to a certain degree. As noted in the introduction, the epistemological tools used to analyse group beliefs will have to be adopted to accommodate the group case. Which adoptions are necessary or advisable depends on the account used to explain what a group, in non-summative terms, is understood to be. In this sense, epistemology depends on ontology. On the other hand, from an epistemologist's perspective, which group account is chosen will partly depend on which is best suited to the epistemological endeavour. From a normative perspective, such as epistemology's, a social ontological underpinning will depend on the epistemological framework. In the course of this thesis, further ways in which social ontology and epistemology depend on one another will indeed be encountered.

To accommodate this relationship and to narrow down the phenomena discussed in this thesis, it will be advisable to restrict both fields. This way a workable framework for further discussions will be established. In Chapter 1.1, I narrow down the very kinds of group that this thesis is concerned with, namely those that qualify as non-summative. Non-summative assumes that groups are, or at least can be, entities in their own right. More precisely, the kind of non-summative assumed in this thesis necessarily involves some intentional or active part on behalf of its members. This constitutes the social ontological assumption. While this assumption should allow for various non-summative accounts, this thesis will mainly discuss List and Pettit's model of group agency (see Chapters 3 to 6).

¹¹ Still, it might be that for an epistemologist this thesis privileges ontological questions; while for ontologists this thesis privileges genuine epistemological matters. I consider this project successful if both parties are indeed uncomfortable with its scope.

Chapter 1.2 is concerned with the epistemological framework and introduces an externalist epistemological framework that ranks among the standard approaches of current epistemology, namely process reliabilism. This constitutes the epistemological assumption that guides the discussions undertaken in the thesis. Intuitively, externalism appears to be the better choice when analysing group knowledge,¹² at least given non-summativism, as introduced in Chapter 1.1. The reason for this is that externalism, as opposed to internalism, does not require the subject to have any access to the very processes that render a belief justified. This is an advantage since groups are likely to be incapable of relating to their own attitudes in ways individual human beings are.¹³

Finally, Chapter 1.3 elaborates upon a working hypothesis for what might present a central challenge to group epistemology given non-summativism about groups and a process reliabilist framework. This enables a determination to be made regarding which aspects about groups and group beliefs are particularly relevant to becoming clear on justified group beliefs and knowledge, thereby guiding the focus for later discussions.

1.1. Non-Summativism about Groups

We often hear that companies, associations, parties, countries, courts or other institutions want, believe or intend something. A quick look at *The New York Times* and a selection of its headlines clearly demonstrates this: “Utah Creates a Registry for White-Collar Crime”, “French Court Rules it Has Jurisdiction over Facebook in Nude Painting Case”, “ISIS is Adept on Twitter” or “Bento Wants to Be Your App Matchmaker”.¹⁴ Non-summativism claims that such talk can be more than just a shortcut for the “people in Utah”, “judges sitting in the court in France” or “people who seek to implement a certain ideology by means of violence”, for instance. Rather, non-summativism holds that when talking of courts, companies, states and other groups it might be apt, although need not be necessary, to depict certain groups as autonomous entities and hence as existing over and above their members.

¹² Pace Hakli (2007).

¹³ However, I do introduce a group account in Chapter 7.2 in which a group relating to its own attitudes is not only possible but also rationally demanded.

¹⁴ New York Times, online 12.3.15

Yet such potential entities are not the only cases in which we talk about groups. Even the example provided herein may not describe the same types of group. In order to become clear on the idea of non-summativism about groups – the very idea of a group discussed in this thesis – allow me to explain the specific kinds of group that this thesis is concerned with by means of an example. What makes a group a group in the sense relevant herein, as opposed to a mere collection of people?

An initial response might state that we must presuppose that a group is a collection of two or more individuals that have a particular feature in common. If so, everyone walking along Universitätsstraße at any one moment may be considered part of the same group. Yet apparently this is not what we mean when talking about group accounts in the sense relevant to this or most probably any other project concerning itself with groups and group agency. Some further or indeed different condition is needed. A common suggestion is that there must be some goal or interest that the members of a group have in common. Or alternatively, another option would be to extract a certain framework of shared experience from a collection of people that allows for us to refer to this collection as a group. For this latter approach to make sense there would have to be a more specific shared experience than merely walking along Universitätsstraße. For example, it would be sensible to argue that some human beings are ascribed the status of “women” and that, because of this, these human beings experience a certain systematic treatment that differentiates them from how “men” are treated, particularly in terms of restrictions or behavioural norms. This represents a shared framework of experiences, which in turn constitutes or entails the identity associated with “being a woman” or “being a man”, and which these collections of human beings share. This allows us to refer to those human beings that share such an identity as a group. Indeed, such an approach is of particular interest to, for example, sociology and feminist philosophy, philosophy of race and other disciplines both within and outside philosophy. I do not wish to diminish the significance or relevance of such approaches in any sense. However, in the course of this thesis a criterion by which to mark non-summativism about groups is sought after and since such

groups can most likely be reduced to their members by means of the ascriptions made to them, it is clear that we are looking for different phenomena.¹⁵

Let us return to the people walking along Universitätsstraße at any one moment. Any selection of people walking along Universitätsstraße at any one moment does not constitute a group simply by means of their walking along Universitätsstraße at the same time and thus sharing a comparable experience. Something else is needed to mark the aspect that makes the difference. One suggestion might be as follows. Not every person walking along Universitätsstraße at any one moment constitutes a group, however, every person on Universitätsstraße at any one moment making their way to a particular coffee shop does, for example. In this scenario, those walking along Universitätsstraße at any one moment who share the same goal or intention may form a group.

However, when talking about groups that qualify as potential agents, it seems that we do not have such groups in mind either. The mere fact of possessing overlapping goals does not suffice. Just because I want to buy a coffee at coffee shop x and you want to buy a coffee at coffee shop x does not simply allow for us to claim that our belief as to where the coffee shop is situated is an instance of a collective belief, as we shall see in the following section in this chapter. Initially, it would appear rather implausible to state that a group would somehow result from the collection of people walking towards the same coffee shop, which is somehow independent of the said people and forms the belief that the coffee shop is in Berggasse. Another suggestion might be that the members of a group not only have to share some goal in common but that, in addition, they have to be aware of this fact in order for a group to emerge. In such a case, a number of people queuing in the coffee shop and waiting to buy a hot beverage may qualify as a group to the extent that the individuals queuing are aware of each other's intentions.

Does this characterisation suffice? The answer still appears to be “no”, although the distinction may not always be that easy to determine.¹⁶ The important thing to show at

¹⁵ Note that I do not wish to exclude the idea that analysing such phenomena may include philosophically interesting instances of collective knowledge. There is nothing wrong with the idea that there are different kinds of collectively laden knowledge. My concern in this thesis, however, is with specific kinds of group and hence with knowledge that emerges in these kinds of group. In the case of the example about women, the groups I am interested in here, as we shall see in a moment, are those that emerge when some women are organisers, for example, in a women's rights movement.

this point is merely that there exists a sense in which the people sitting in a bus are not (necessarily) a group in the sense that is relevant here, while a research group publishing a paper together most likely is. Another example that can be used to mark the difference is as follows. Imagine five people standing in an elevator in order to go to the upper floors of a building.¹⁷ These people certainly comprise a collection of individuals and we may call them a “group”, albeit in a different sense from that which interests us here. However, imagine that the elevator has a malfunction and stops suddenly. At the moment that these people try, let us assume, to organize their efforts in order to get out or call for help something changes. While they were a mere collection of people at the beginning, they now find themselves engaged in a joint effort to get out of the elevator. This kind of group is closer to the kind of phenomenon that this thesis is concerned with. Joint intentions and commitments (Gilbert 1990; 1992), shared intentions (Bratman 1999), intentional inter-connections (Schmid 2008), we-intentions (Searle 2002), we-mode (Tuomela 2010), joint rational viewpoints (Rovane 1998) and so on. All of these criteria and more have been used to mark the aspect that makes the difference.

Returning to the coffee shop example, whatever the difference may be, even if I am aware that the people in front of me are also queuing to buy a hot beverage, we do not constitute a group in the sense that a number of robbers might if they were queuing in order to wait for the right moment to rob the coffee shop together.¹⁸ I might very well believe that the person in front of me is queuing to buy some hot beverage, and she might believe that this is why I am standing next to her too; and if someone tries to pass through the queue that we are standing in I might step back and she might step forward, thus creating some space in order to let the person pass. In this case, it may or may not be so that we act jointly, and it is even more questionable as to whether we constitute a group. But at least it would appear to be an over interpretation to say that in such an

¹⁶ This is difficult to determine because the behaviour of this group might not be distinguishable from the behaviour of a group. See, for example, Searle (2002) and see also Schmid (2012) for a related discussion.

¹⁷ I owe this example to Felix Brodbeck.

¹⁸ Note that Schmid uses a comparable example, which also includes a robbery (see Schmid 2012, esp. p. 55f). His aim is to show that rather than mere behaviour, collective intentionality is key when analysing matters of collectivity. Yet groups may or may not exist in a special form of collectivity (depending on the chosen account), and they may or may not necessarily involve collective intentionality, as will become clear shortly. Hence my example probably aims at something slightly different.

instance a group exists in its own right of which the person standing in front of me and myself are both members, and that this group decided that we should move and so we, as its members, did. In the case of the robbers, however, things look different. Their queuing has been coordinated and planned and they each intend to execute their role in the robbery. They will have planned this for a certain time span and their standing in the queue is most certainly an action that is jointly performed. While in one context or another all of the aforementioned collections of people may rightly be referred to as a “group”, the cases of the robbers or the people caught in an elevator are of interest to us here.

Yet while the above phenomena come close to explaining non-summativism about groups, there exist various explanations that address the ontological status of such collections of people that need not lead to the emergence of a collective entity. For instance, John Searle (2002) would claim that the robbers or the people caught in the elevator are linked in the sense that their attitudes are of a collective rather than an individual nature. They jointly perform actions and form joint intentions in that each person adopts a “we” rather than an “I” attitude. Consequently, not each and every person seeks to get out of the elevator but instead “we” seek to get out of the elevator. Equally, not each and every one of the robbers seeks to rob the coffee shop but “we, the robbers, want to rob the coffee shop”. However, in such a scenario a group mind remains absent and a separate entity fails to emerge since these joint attitudes are merely a psychological state or mode of intentionality in the individual minds.¹⁹ Yet for Margaret Gilbert the contrary holds. According to her plural subject theory, a plural entity does indeed emerge from all instances in which we form joint intentions or act jointly (see Gilbert 1990; Gilbert 1992). Hence her account is clearly non-summativist in the sense that is relevant to our discussion. List and Pettit’s story varies significantly from both Searle and Gilbert’s approach, despite the fact that like Gilbert’s it too is non-summativist. They would certainly wish to keep the example of the robbers apart from that of the people caught in the elevator. (See List and Pettit 2011) The former might qualify as members of a group agent that exists in its own right, while the latter simply comprise a collection of people

¹⁹ Searle’s account of groups approaches the issue from a different angle in that groups are a social fact that derives from imposing the status function on a selection of people in that they count as group G in (social) context C (cf. Searle 1995).

involved in a shared action along Bratmanian lines (see *ibid* and Bratman 1999). Further still, if we turn to Hindriks (Hindriks 2014) yet another story can be told. If the people caught in the elevator collectively accepted some decision mechanism, they would constitute a group agent. In this account the requirements are less restrictive than in List and Pettit's model, but they are probably more demanding than in Gilbert's. Quite distinct from all of these approaches is Carol Rovane's idea of joint activity (see Rovane 1998), in that for her, the aim of achieving overall rational unity on the collective level constitutes group persons. How, then, are we to pinpoint non-summativism?

Luckily, this discussion requires us to do so in a minimal sense only. For this thesis the assumption suffices that under certain conditions plural, corporate, or group entities can emerge and that those entities are somehow autonomous with respect to their constituents, that is, the individuals comprising it. Any account that includes such an idea will potentially qualify for discussion in this thesis.²⁰ Still, for the most part this thesis is restricted to a discussion of one particular account, namely List and Pettit's group agents.

What, then, constitutes non-summativism about groups in the sense relevant for the discussion in this thesis? As we have seen, mere instances of shared experiences (walking along Universitätsstraße), of common knowledge (knowing that others are also queuing to buy a hot beverage), and of simple "coordination" (two people each letting someone pass through the queue) may all rightly allow for the labelling of a collection of people as a group, but they do not constitute the kinds of group that non-summativism is interested in. Nor do accounts of collective intentionality necessarily lead to non-summativism. What needs to be in place is merely that there exists something that makes it the case that a group emerges or is sustained as a separate entity. How exactly this autonomy of a group with respect to its members comes about and/or what exactly this "glue" consists of is not (yet) of relevance and may indeed be the result of some joint attitude; some active identification process with the group; a model of interrelated intentions; merely the existence of some organizational structure that implements the

²⁰ Hence among contemporary theories this includes, for instance, the following accounts: Margaret Gilbert's plural subjects, List and Pettit's group agents, Carol Rovane's collective persons, Schmid's account of plural actions. Also note that in the following I present further restrictions that go beyond non-summativism.

aggregations of the judgements of the members; reciprocal commitments and obligations; some kind of “we-mode” and “ethos” (Tuomela 2010); stability over time, agreed upon decision mechanism or what have you. I take it though that accounts allowing for autonomous collective entities share at least some features. Accordingly, these features, which are probably necessary but most certainly not jointly sufficient to denote the very accounts discussed here, could be the following:

- (I) There are groups that exist as independent entities to which it makes sense to ascribe certain attitudes and features.

Without doubt, condition (I) presents the core idea of non-summativism. This idea is that groups are not merely the summation of their members and that instead they can be something over and above that. However, non-summativism alone does not necessarily state that corporate attitudes also are irreducible to member attitudes.²¹ This thesis presupposes an even stronger claim and thus a further condition that describes non-reductionism must be added. It might look like this:

- (II) The features or attitudes of such a group cannot be reduced to the attitudes or features of its members.

Still, this picture allows for too many varying group accounts. For example, Bird (2014) claims that analogous to distributive cognition, a singular group attitude can emerge without the members *actively* participating in its development. While I do not wish to exclude this as a possibility, I believe that collective epistemology regarding such phenomena looks quite different from the focus of this thesis. Groups, as understood herein, depend on the active participation of their members in one way or another.²² Therefore, albeit for reasons of convenience alone, a further restriction is useful here in order to exclude those scenarios that might become mixed up with the kinds of group that this thesis is concerned with.

²¹ List and Pettit’s discussion, which distinguishes between realism/non-realism and redundancy/non-redundancy, probably works quite similarly to the distinction drawn here. Moreover, Fagan (2012) seems to make a comparable point. For an example of non-summativism (“realism”) and reductionism (“redundancy”), see Sylvan’s (2012). A related discussion regarding the distinction drawn here can be found in Pettit (1993).

²² Albeit perhaps only loosely, as in List and Pettit’s case.

(III) Such groups include members who play an active part and who are (as thus) intentionally involved in constituting or sustaining the existence of the group, as well as its attitudes and features.

Note that intentional involvement need not necessarily mean that such intentionality is collective in nature.²³ But it does not exclude that these intentional involvements are or perhaps even have to be of a collective nature either.

In light of the above, what I refer to in general as “non-summativism” actually amounts to a threefold claim (I)-(III). Taken that group entities are able to emerge as something independent of their members, this is non-summativist in a narrow sense. In addition however, non-reductionism about collective attitudes is also included here and the members of a group have to be intentionally involved, somehow, in group matters. Yet since the most essential claim from this threefold structure is (I), I refer to these three restrictions as the non-summativist assumption.

1.2. Epistemic Externalism and Process Reliabilism

In the above section, the types of group accounts discussed in this thesis have been defined. I label these accounts as simply non-summativist, although non-reductionism and intentional involvement on the part of members are also assumed. It is now time to turn to the second presupposition, namely the epistemological one.

When investigating the possibility of group knowledge, we must first be clear on what qualifies as knowledge as opposed to mere belief. It has long been assumed that two further ingredients, despite that we are dealing with belief states, are needed here. First, in order for some belief in *p* to count as knowledge, a belief needs to be true. Second, truth alone is not enough to constitute knowledge; a belief needs to be justified as well. You might believe that a secret alien invasion is taking place and rightly be accused of partaking in a conspiracy theory. Your “reasons” for this belief merely are that a strong emotion triggered out of the blue which somehow makes you believe that there is an alien

²³ That is to say, holding attitudes in first person plural terms, instead of first person singular terms. For instance, “we want to *x*” instead of “I want to *x*”. See for example Bratman (1999) for an account in singular terms and Searle (2002) for “we-intentions” as primitive phenomena.

invasion going on.²⁴ But suppose that we are in some far away possible world in which it is indeed the case that a secret alien invasion is taking place. Let us further assume that your evidential situation is the same, that is to say, you do not possess any sufficient reason to determine the issue. Your belief would be true, but would we really attest that you know that there is a secret alien invasion going on? We certainly would not. But why not, given that your belief is true? The answer is that despite your belief being true, something is at odds with the justification you may (or may not) believe to have and in fact do not have. We therefore see that justification takes centre stage in providing the answer to the question regarding what it takes to know a proposition rather than merely believe it. Hence to know that *p*, we need a justified true belief (JTB).

Since the publication of Edmund Gettier's *Is Justified True Belief Knowledge?* (Gettier 1963), it has become clear that the JTB conditions for knowledge may either be incomplete or that the justification-clause must be fundamentally revised. As Gettier shows, the problem lies in the fact that we are able to construct cases in which a subject believes a true proposition, *p*; the belief is justified, but the subject does not know that *p*. A simple Gettier case can be constructed as follows. A simple Gettier case can be constructed as follows.²⁵ Imagine you are looking at your watch and it tells you that it is 9 o'clock. Indeed, it is 9 o'clock. Hence your belief in the proposition *it is 9 o'clock* is true. It is further justified because you were looking at your watch and thus had sufficient reason to believe that *p*. However, unbeknown to you, your clock stopped several hours ago. As chance would have it, you looked at your watch precisely twelve hours later, that is to say, at one of two possible times a day at which your watch indicates the correct time. Hence your belief is justified according to JTB and it is true, yet many epistemologists hold that, intuitively, you do not know the time.

One reaction to cases like this has been to fundamentally revise what justification might amount to. While in traditional epistemology justification amounts to the reasons that a subject has, or at least in principle has access to an alternative strand, namely externalism, denies this requirement. Externalism shifts the perspective away from the reasons that a subject may or may not consciously have for believing that *p*, and

²⁴ I am grateful to Sebastian Kletzl for making helpful suggestions as to how to improve this example.

²⁵ This example was probably first used by Russell (see Heathcote 2012).

alternatively moves towards asking what has to be in place for a particular belief to qualify as knowledge, independent of the subject's own assessment of the matter.²⁶ Externalism therefore emerges as the denial of traditional internalistic perspectives regarding justification. Not only is externalism a reaction to Gettier cases, but it is also a reaction to several problems that internalistic accounts of epistemic justification traditionally face. One such problem runs as follows. Internalism classically assumes that a belief is justified if the *subject has* good epistemic reasons for her belief. The idea is either that these reasons are, in fact, available to the subject, or that they can, at least in principle, be made accessible to the subject. In assuming this, however, we run into a regress problem (see, for example, Bonjour 1980[2001]). Given that reasons are available to us in the form of beliefs about certain facts, if we need good reasons to not only believe but also justifiedly believe that *p*, then we further need justification for those reasons. After all, what makes them “good”, in epistemic terms, is that they are justified. This constitutes a regress with respect to justification. In order to be justified we must be justified in those beliefs that are supposed to justify our beliefs. An attempt to stop this regress traditionally turns to foundationalism. The assumption here is that some beliefs are not in need of further justification. We have “basic beliefs” (ibid), like those that stem from perception, which are justified by default and which then serve as the basis for other beliefs, which then provide reason for yet other beliefs to be justified and so on and so forth.

Yet even if we assume that foundationalism has a point, we still run into problems. If it is a requirement of justification that for a belief to count as justified it must indeed be just that in the believer's eyes, or at least must thus in principle be accessible to the believer, she will need reasons to refer to some of her beliefs as basic. However, these beliefs about her basic beliefs will not be basic and as thus are once again in need of justification. Accordingly, we are again stuck with a comparable problem. (See (Bonjour 1980[2001]))²⁷

As already mentioned, an externalist perspective is characterized by dropping the requirement that justification must be a state that the subject has access to in principle, or perhaps even has to consciously find herself engaged in. Thus it also opposes the idea that

²⁶ See, for example, Kornblith (2001) and Bonjour and Sosa (2003) for an overview.

²⁷ I will not go into detail here since internalism need only be discussed to the extent needed to introduce epistemic externalism and in particular process reliabilism.

epistemic justification is bound to a “regulative” (Goldman 1980[2001]) function that should enable a subject to evaluate or determine the justificatory status of her belief. Rather, ascriptions of justification and knowledge are independent of the subject’s own evaluative attitudes regarding her beliefs and, furthermore, must not be accessible by her in principle (cf. *ibid*). In this sense, externalist theories evaluate a belief-state from “outside” the believer. Externalism usually considers “an *already formed* belief of a cognizer” and, given this belief, asks “what features are necessary and sufficient for that belief to count as justified” (*ibid*, p. 38)? Given such a perspective, the regress problems mentioned above can be avoided.

The most prominent position within standard reliabilist epistemology is probably Alvin Goldman’s process reliabilism. Process reliabilism essentially holds that a belief is justified in the case it has been formed via processes that are reliable, where a process qualifies as reliable if it produces more true than false beliefs. Such processes should thus be truth-conducive. However, in order to avoid question begging they should not be of an evaluative kind or involve “justification” in conceptual terms. Accordingly, reliability should be explained using “non-epistemic terms”²⁸ (Goldman 1979[2012], p. 30, and Goldman 2012a). For instance, while in the course of rational belief formation certain processes may very well constitute the “non-epistemic” processes essential for reliability, it cannot be a requirement *that belief formation be rational* because rationality as a concept is too close to a concept of justification. Truth does not itself entail justification and is therefore one factor that process reliabilism can and does consider a necessary condition for knowledge. Furthermore, the neuronal processes and other mainly causal factors involved in belief-formation, such as cognitive and psychological states like “visual and memory experiences” (Goldman 2012a, p. 72), count as such conditions. Accordingly, the neuronal, psychological and causal histories of a belief’s occurrence, if reliable, justify that resultant belief. If a belief not only possesses a history of reliable processes that led to its occurrence but is also true, then that belief qualifies as knowledge. This focus on the processes that precede a belief explains why Goldman repeatedly refers to his process

²⁸ This does not exclude the idea that beliefs often are based on other beliefs. However, the idea is that in each case beliefs will be justified if they are reliably formed. This does not exclude that the reliability of one belief determines the reliability of another. Also note that the varying ways in which to phrase reliabilism are not mutually exclusive. Most notably, Alvin Goldman’s process reliabilism has ecumenical tendencies (see, for example, Goldman 2011).

reliabilism as “historical reliabilism” (Goldman 2012). It is a historical theory that essentially diverges from classical epistemological approaches due to the following aspects of justification:

A reliable inference process confers justification to an output belief, for example, only if its input beliefs were themselves justified. How could their justifiedness have arisen? By having been caused by earlier reliable processes. This chain must ultimately terminate in reliable processes having only non-doxastic inputs, such as perceptual inputs. Thus, justifiedness is often a matter of a history of personal cognitive processes. This historical nature of justifiedness implied by process reliabilism contrasts sharply with traditional theories like foundationalism and coherentism, which are “current time-slice” theories. (ibid., p. 73)

There exists a chain of reliable processes that underlies both inferential and non-inferential beliefs. Focussing on this chain diverges from the idea that justification can be determined given the current state of affairs that the subject is in at time *t* at which she believes that *p*, because it stresses the “history” of the belief’s formation. Since this constitutes an externalist account, it is not required that the subject actually be aware of her belief’s reliable or unreliable history of formation or its truth.

Of course Goldman’s suggestion has been the target of many objections,²⁹ the result of which the theory has gone through many, partly fundamental, adaptations and amendments. Some of these refinements concern the criteria required to meet certain possible world scenarios in which the underlying processes of a belief fail to be reliable or fail to appropriately handle various evil demon arguments. Goldman’s response has been to add a clause to process reliabilism that states that reliability must only be stable over “natural possible worlds” (Goldman 2012a).³⁰ Another adjustment has been made with respect to the extent to which, or indeed the ways in which, processes count as reliable. In addition to processes being “globally” (or unconditionally) reliable, which means that they are in any case truth-conducive, independent of the very situation or input, Goldman further adds the possibility of local (or conditional) reliability, in which a certain process

²⁹ See Goldman (2012b) for an overview.

³⁰ He thus diverges from Nozick’s tracking account in that Goldman seeks to completely shut out any scepticist argumentation, while Nozick’s account takes scepticism seriously and allows for scepticist arguments to at least partly make a point. (Goldman 1983; Nozick 1981)

is reliable depending on other factors.³¹ Despite turning away from internalism, Goldman has made further suggestions regarding how process reliabilism can help a subject when evaluating her own beliefs, as well as how certain internalistic theories may potentially be subsumed in his framework (see Goldman 2011). Taken as it stands, one particular worry might be that the complete reliability of all the processes that lead to a certain belief is a too strong a requirement for justification. Accordingly, suggestions have been made with respect to the degree to which a belief's history must be reliable. In particular, Goldman argues that the reliability of a set of processes that precedes a certain belief need not be reliable in total, but can only be so "locally".³²

Not all of the refinements suggested by Goldman himself have survived further scrutiny and some are yet to undergo revision.³³ The process of refining and complementing the general idea of process reliabilism into a more detailed theory is hence not yet complete, as we shall see in Chapters 6 and 7, in which Goldman's very recent proposal of "social process reliabilism", that is, process reliabilism applied to group beliefs, is discussed (cf. Goldman 1999, 2010, 2014).

Moreover, in addition to Goldman's own refinements, there are various versions of process reliabilism that diverge at significant points from Goldman's own account, but which agree with the general idea of justification being phrased in terms of the reliability of a belief's underlying (causal) processes. Examples of such theories include Ernest Sosa's and John Greco's (quite different) virtue epistemologies, as well as Sanford Goldberg's non-reductive account of testimony and general anti-process individualistic process reliabilism.³⁴

As we have seen in this section, this thesis is committed to a certain ontological framework regarding groups. Here non-summativism provides a narrow framework of relevant group accounts, but it does not apply to one specific theory alone. Equally, this

³¹ An example of local reliability would be inference. In the case of inference, the reliability of the inferential process depends on the reliability of the input beliefs from which one infers something (cf. for example, Goldman 1986).

³² Imagine an instance of overdetermination in which a person believes that *p* based on a great number of different processes indicating *p*. Among these many processes, there is only one process that is slightly flawed. It would be too high a requirement to claim that, because of this one flawed element, the belief has been unreliably formed.

³³ Again Goldman (2012b) provides a comprehensive overview of this topic.

³⁴ See for example Sosa (2009); Greco and Turri (2013); Greco (2010); Goldberg (2008 and 2010).

thesis is committed to process reliabilism yet remains open in terms of which version will prove the most useful in our discussion of the possibility of epistemically justified group beliefs. However, since Goldman's account is perhaps the most notable account in current existence, presents the first discussion of how to apply process reliabilism to group phenomena, and has served as the basis for other process reliabilist accounts, a brief discussion of his version of process reliabilism constitutes a natural starting point. The basic features of process reliabilism noted above can be roughly expanded to relate to other process reliabilist accounts other than Goldman's.

Before turning to the next chapter a further note is in order, this time with regard to the simple matter of vocabulary-choice. Above I use the phrases "group knowledge", or "justified group belief" and its simpler version "group belief". As we have seen, according to process reliabilism, the difference between these justified belief and knowledge merely consists of whether or not, in addition to a belief's reliability and hence justification, the respective belief is also true. The epistemological challenge lies in the justification-clause and while at certain points in my discussion a belief's truth in addition to its justification will turn out to be important, this is more often not the case. Therefore, due to convenience I sometimes use the phrases "group knowledge" and "justified group belief" interchangeably, in particular where it remains unimportant as to whether or not, in addition to justification, we also assume a belief to be true. Furthermore, beliefs and (more generally) attitudes can be justified in various ways. They can be morally justified or rationally justified to the extent that beliefs depict the right means to an end. In this thesis, the term "justification" is used in reference to epistemic justification, unless otherwise indicated.³⁵ Further, unless explicitly stated otherwise, I use the term "justified" in a manner roughly synonymous with "reliable" or "reliably produced", as is customary in the respective process reliabilist literature.

Hence, these are the two preliminary assumptions that will guide this thesis' discussion. On the social ontological side I assumed non-summativism, understood as a threefold

³⁵ Note, however, that in Bayesian epistemology, for instance, epistemic justification and "practical", that is to say, action-oriented justification are sometimes seen as the same thing to the extent that in each case our beliefs are rational, and hence justified, only if they were formed in accordance with the laws of probability.

claim including not only the idea that there exist groups in their own right, but also that a group's attitudes are irreducible to those of its members. Moreover, for reasons of convenience, I further assume that a group's members need to be somehow intentionally involved in a group's emergence or sustainment. On side of the epistemological framework this thesis is committed to externalist theories regarding knowledge and in particular to versions of process reliabilism. In the next part (chapters 3 and 4), one representative of the social ontological assumption is introduced. List and Pettit's account of group agency, which has served as a basis for process reliabilist discussions regarding group belief is discussed in greater detail. In part 3 of this thesis, I turn to a process reliabilist discussion of justified group beliefs on the basis of the ontological backdrop of List and Pettit's group agents.

Chapter 2

Towards an Epistemology of Groups. Some Essentials

With the two preliminary positions that this thesis adopts to hand, it is now important to clarify some general issues that stem from the discussion so far and relate to an examination of justified group belief. There are, in particular, two issues that need to be addressed that will help to determine the direction in which we should turn if we are to become clear about the possibility of justified group belief. First, more needs to be said about the relationship between group members and a group. Non-summativism is bound to the claim that groups can exist in their own right, yet if there are no members then there is no group. After all, the members constitute the group. However, this appears to exist in a state of tension with any claim of autonomy that non-summativists might make. The extent to which – and ways in which – they can identify autonomy is defined and discussed in the course of introducing List and Pettit’s group agents. Yet at this point in the thesis it is useful to draw a general, albeit sketchy picture of the dependency that exists between members and a group. This is also essential for delineating – at this point, in an equally sketchy manner – where our discussion regarding justified group belief should commence.

Second, it is now useful to identify some features that must be applied to collective instances of justified belief. This way, it should also be possible to shape what might be the most important question concerning how groups may or may not be capable of holding epistemically justified beliefs. As previously mentioned, our epistemic tools are made to measure for individuals. Groups are not usually examined in standard epistemology and hence a certain flexibility in adjusting our epistemic tool set is required. Accordingly, a first approximation of what adjustments probably need to be made and what challenges group epistemology will probably have to face is in order. All these assumptions are merely working hypotheses and, albeit loosely, bear on existing discussions about justified group belief.³⁶

³⁶ In particular, however, see Mathiesen (2006).

This thesis primarily discusses one non-summative account that, together with Margaret Gilbert’s plural subject theory, dominates both current social ontological and existing collective epistemological debates. This account is List and Pettit’s conception of group agency as developed primarily in their *Group Agency* (List and Pettit 2011) as well as a series of papers published jointly as well as individually and in cooperation with other authors.³⁷ I believe that most non-summative group accounts are compatible with the picture illustrated in Figure 1, which essentially displays the relationship between members and their group in non-summative as introduced in the first part of the above chapter.³⁸ We could determine that these arrows display the ways in which, generally speaking, members and groups determine each other. Alternatively, these arrows might indicate relationships of responsibility or they might depict relations that concern more constitutive relationships between members and a group. However, since our concern lies with epistemically justified *beliefs* then the arrows should instead indicate that the attitudes of the members that determine that of the group and vice versa. In particular, while List and Pettit seem to focus on the upper arrow, thus asking how a group can come to hold a certain attitude given its members’ attitudes, Gilbert is, for example, instead concerned with the lower arrow.




Figure 1 Relationship between Members and a Group

³⁷ These publications include, among others, Dietrich and List (2004, 2013); List (2005, 2006, and 2014); List and Pettit (2002, 2004 and 2006); Pettit (1993, 2006, 2007, 2010 and 2014); Pettit and Schweikard (2006).

³⁸ Note that the picture might be applied to various aspects of non-summative. We could determine these arrows as displaying the ways in which, generally, members and groups determine each other, or these arrows might indicate relationships of responsibility, or alternatively they may depict perhaps even the relations concerning which constitutes what. However, since our concern lies with epistemically justified beliefs the arrows should rather indicate the ways in which the attitudes of the members determine that of the group and vice versa.

As simple as Figure 1 may seem, it does depict a central issue in discussing the possibility of justified group belief. Since groups are comprised of members, it must be the members who somehow constitute a group's set of attitudes. This is what the upper arrow should indicate. Both of the aforementioned accounts allow, in principle, for cases in which the group holds a certain attitude, even though not a single member of the group holds that same attitude; that is to say, it might be the case that a group believes that p , while all its members believe that $\neg p$. At first glance, this puts into question the relevance of the upper arrow. Since this concern has been discussed at some length in recent debates (in particular with respect to Gilbertian plural subjects), and since List and Pettit also claim that group attitudes can radically come apart from member attitudes, a few words regarding this particular problem are in order.³⁹ This will allow us to anticipate important concerns within the discussion presented in this thesis in the chapters that follow.

There are many ways in which extreme divergences between group and member beliefs can come about. Leaving out the technicalities, some ad-hoc possibilities serve to explain this situation. One scenario might be that the group always adopts the opposite of its members' majority beliefs. Another less absurd option might be that a group decides to let an external expert, or expert panel, decide upon an issue. Yet another example may include a case in which a group belief formed at t_1 conflicts with updated beliefs among the members at a later time, t_2 , based on, for example, new evidence that is not (yet) available to the group.

I mention these extreme cases because they allow us to spot what has been identified as the crunch point of collective epistemology concerned with knowledge in non-summative group accounts (cf. Goldman 2014; Mathiesen 2006). In the course of the discussions herein, it will become clearer as to how such instances might occur. In the last chapter, moreover, we will see that in certain non-summative accounts, like Rovian group persons, such instances are impossible. For our current purposes however, the

³⁹ Note that in an epistemological context, Gilbert's plural subject account has been mainly discussed with respect to the lower arrow displayed in Figure 1. See, for example, Gilbert (1987 and 2004); Wray (2001); Mathiesen (2006).

important thing to note is the following. Even if group beliefs and member beliefs diverge, this does not mean that a group's beliefs (or attitudes more generally) do not hinge in some essential way on member beliefs or attitudes. In fact, such cases rather point out that there is an epistemological question to be asked prior to dealing with such extreme cases, namely how can a group come to believe that *p* in the first place? Moreover, how can it justifiably come to hold a belief?

The central challenge of this thesis is to give an account of how a group is able to come to believe something in a justified way based on its members' roles in bringing about this belief. This enables us to make sense of various extreme cases like those mentioned above. For example, it is demanded from a sound epistemological account of justified group belief to be able to deal with the abovementioned ways in which a group's belief can come apart from its members' beliefs. It should also be able to determine which of these options warrant justification on a group level and which do not. For instance, in the first scenario we would most likely want to attest that the group belief is unjustified (except, perhaps, if the group comprises members who are very likely to form false rather than true beliefs). In the second scenario we would want to attest that the group belief is justified, given that the expert provides a reliable belief and the members are bad judges as to whether *p*. The third scenario would prove more complicated to determine and probably depend on a more fine-grained picture of the given scenario.

Be that as it may, the important thing to note is that, despite extreme cases, in one way or another group attitudes hinge on the members and their attitudes, and thus the relationship depicted by the upper arrow in Figure 1; this will prove the crunch point for discussing the possibility of a group believing something justifiably.⁴⁰

The second lower arrow in Figure 1 that goes from the group and its attitudes to its members depicts the various ways in which group attitudes determine what a member ought to believe or do. One obvious way in which a group influences its members is that a group cannot act on its own. It needs its members to do things for it. Moreover, there are various, and partly epistemically relevant ways in which a group determines its

⁴⁰ It does constitute the crunch point even with respect to Rovian group persons, who are discussed in greater detail in Chapter 7. The reason is that discussing this very point will lead to suggesting a completely different account, such as Rovian's group persons in the first place. Hence, at least it appears to be a beneficial starting point.

members' attitudes. For example, given that a member does not believe that p or is indifferent about p , what is she to do when learning from the group that p ? Initially, an obvious way to make sense of this relationship is to analyse these types of situation as ones in which a member is confronted with a testimonial belief provided by her group. However, it is not initially clear that testimony indeed provides the most proper framework with which to deal with such situations given the unique, though as of yet unclarified relationship between members and their groups. After all, this might be one of the cases where adjusting our tool set might be advisable.

Essentially, we are confronted with two ways in which the relationships between members and groups are important (at least in accounts where drawing this distinction makes sense). First, as depicted in the upper arrow in Figure 1, there exist various ways in which the members somehow determine what is happening on the group level. Second, there exist various ways in which the group has a grip on its members, as depicted by the second lower arrow in Figure 1. While this second lower arrow depicts various epistemically relevant questions, it is this first upper arrow that is of primary import to this thesis. After all, the second arrow is more concerned with what it means for the members of a group to be confronted with a group attitude *once this attitude is in place*. When asking how a group can justifiedly believe something, we are asking – in accordance with the chosen externalist framework – for the very story to be told of how this belief came into existence. Therefore, our focus will clearly lie with the first upper arrow in Figure 1.

However, it is important to keep in mind the fact that the formation of a group belief may not be a one-way street as this upper arrow indicates. Consequently, we need to be clear on the ways in which this lower arrow may play into the story of how a group can come to form (or possess) a justified belief and perhaps even knowledge. Moreover, in Chapter 7, a yet different non-summative account of groups, namely Carol Rovane's idea of group persons, is discussed. It is likely that Figure 1 does not, or at least in a quite different way, apply to Rovane's idea of groups. Yet Figure 1 should not be mistaken for a further restriction with respect to non-summative, but instead as yet another tool to form a first working hypothesis regarding where the crux of the matter might lie.

To sum up, the key to further discussions regarding justified group belief will be to clarify the very epistemic relations between member beliefs and a group belief. Moreover, the ways in which a group's belief formation depends on the member beliefs will be of particular importance. Before dealing with these matters, however, it is important to elucidate the particular framework within which such a discussion will take place. Therefore, a detailed discussion of List and Pettit's account of group agents is in order.

PART II

Chapter 3

List and Pettit's Group Agents

In recent philosophical discussions about groups, List and Pettit's model of group agency has probably gained the most attention.⁴¹ Their account clearly counts as “non-summative” and has been laid out in greater detail in their book *Group Agency* (2011). Their particular version of non-summative is called “non-redundant realism” (NRR) about groups. (See in particular List and Pettit 2011, pp. 2ff) It is this particular account that dominates most of the discussions in this thesis. One reason for this is that it is the only non-summative group account to have been discussed at length by a process reliabilist, namely Alvin Goldman. Accordingly, the next two chapters are dedicated to this particular account.

After introducing their account and the specific ways in which List and Pettit's account qualifies as non-summative in this chapter, Chapter 4 discusses the epistemic implications of their model and how they believe that they meet the “knowledge challenge” in groups. Chapter 5 critically reassesses this discussion and draws some lessons from it. In particular, it identifies one reason as to why List and Pettit's account of how to meet the “knowledge challenge” would be problematic if it were to constitute a fully fledged theory of group knowledge. Furthermore, certain potential problems that emerge for List and Pettit are identified and it is acknowledged that these also constitute a challenge to process reliabilism. In the next part of this thesis, then, I will enter into a discussion of Alvin Goldman's social process reliabilism and discuss the extent to which his approach might be capable of meeting these challenges. As we will see, his account has some problems of its own. In the last chapter of this thesis, I will suggest two solutions to these problems. Both involve concessions towards anti-individualism, which Goldman would most certainly be unwilling to consent to.

⁴¹ See, for example the discussions in the course of a special issue on List and Pettit's group agents in *Episteme* (Briggs 2012; Sylvan 2012; Cariani 2012; Gaus 2012; Dewitt 2012; Dunn 2012; Hazlett 2012), as well as in *Erkenntnis* (Kusch, Pauer-Studer, and Schmid 2014; Kusch 2014; Pauer-Studer 2014; Schmid 2014; Rovane 2014; Hindriks 2014).

List and Pettit identify two kinds of non-summative group accounts, both of which are akin to their own, yet both of which entail certain vices that List and Pettit seek to avoid, while adopting the associated virtues. These accounts are what they call the “authorization” and “animation” accounts of groups (see *ibid.*).⁴² Both of these accounts are non-summative in spirit, but while the first only assumes a thin realism by being mainly redundant with respect to member input, the latter is clearly non-redundant but comes at a price.

The authorization approach can be traced back to authors such as Thomas Hobbes or Jean Jacques Rousseau (Rousseau 2011; Hobbes 1996). The idea is that a collection of people may authorize (and thereby create) a social entity, which speaks and acts on their behalf. In this sense, the resulting social entity comprises members but nonetheless enjoy certain, albeit rather weak degrees of autonomy. According to List and Pettit, the purpose of these accounts is to make sense of human cooperation by giving an account of how best to make sense of “shared voices for certain purposes” (List and Pettit 2011, p. 9)

Exhibiting opposition to these accounts are those that claim that there exist social entities that essentially exceed the collective aspect of the authorization model. List and Pettit label approaches such as these “animation” accounts.⁴³ The key assumption of animation theories is that the emergence and functioning of a social entity is not bound to its members’ authorization, but that it instead constitutes a fully fledged entity in its own right, thereby exercising more or less complete autonomy with respect to its members. Such accounts essentially depend on claiming that social entities somehow have an “essence”, some metaphysical feature that distinguishes them from other, non-collective entities and sets them apart from their members. This “essence” is what gives “life” to such entities and hence warrants their autonomy. (Cf. *ibid.*)

List and Pettit’s account sides with animation accounts (or emergentism) in terms of retaining a strong autonomy claim regarding groups. However, they seek to avoid the vice of having to claim some trans-individual “nature” or “essence” on a metaphysical level. This can be achieved by reference to authorization accounts of social entities. The idea is that the relationship between members and a group is akin to authorization accounts, but

⁴² For a more detailed and similar discussion see Pettit (1993).

⁴³ See Runciman (2008) for an overview and List and Pettit (2011, esp. pp. 8ff).

that this relationship is of a kind that warrants a group's autonomy with respect to its members in a rather strong sense, as assumed in animation theories. As we shall see in what follows, the key to achieving this is claiming that a supervenience relation exists between member attitudes and group attitudes. According to List and Pettit, there exists a specific dependency of group attitudes on those of the group's members, but at the same time they diverge in significant ways from those of their members in that they need not depend on them directly. The assumption is that the group's attitudes supervene set-wise, that is to say, holistically, upon those of its members instead. This supervenience claim is supported by certain "impossibility" results akin to or part of well-known impossibility results in judgement aggregation theories.⁴⁴

All in all, List and Pettit seek a middle way that can be understood as either a strengthened version of authorization accounts or a weakened version of animation theories.

Before turning to the details of List and Pettit's account, one further note is in order. This regards the scope of List and Pettit's accounts of groups. As we have seen in the first chapter, there are various ways to refer to a collection of people as a group. The scope of their account of groups may best be illustrated by comparing it to another. Again, the contrast that makes the most sense regards Gilbert's plural subject theory. Gilbert's account of plural subjects, for example, states that in any instance in which a joint action is performed, a plural subject emerges by means of committing to a joint goal. A plural subject is, at least in certain respects, autonomous with respect to its members. As thus, her account of the plural subject covers many instances and such simple phenomena as two people going for a walk together (see Gilbert 1990 and 1992). This general idea of plural subjects is then applied to conceptions of institutionalized groups. For List and Pettit, such talk constitutes an over interpretation of groups as separate entities (List and Pettit 2011, p. 215f, FN 18). People engaged in some joint action may rightly be referred to as a group, but this does not yet allow for claiming that some separate entity emerges to which we can ascribe certain attitudes. Rather, only certain groups can further qualify as group agents and hence be regarded as groups in a non-summativemanner. In order to

⁴⁴ See, for example Arrow (1963); Bovens and Rabinowicz (2006); Cariani (2011); Dietrich and List (2006, 2009, 2010 and 2012); Knight and Johnson (1994); List and Pettit (2002 and 2004); List (2003, 2004a and 2011); Mongin (2008 and 2012); Nehring (2003); Pigozzi (2006); Sen (1966).

qualify as a group agent and not merely a group, they have to meet certain criteria for agency and rationality. In most cases, the emergence of a group agent will involve a joint intention,⁴⁵ but not just any joint intention. A group agent may come into existence only if the members of a group decide to form a group agent. Therefore, what needs to be involved is a joint intention towards forming a corporate entity, the group. However, such a joint intention alone does not warrant group agency. Given such a scenario, there are certain further criteria that a group has to meet if it is to qualify as an agent in its own right. These criteria are discussed at the beginning of the following chapter.

List and Pettit's NNR makes essentially two claims, one is positive and the other negative. According to the positive claim, *it is possible* to ascribe agency to certain groups. This involves ascribing attitudes to them as well as ascribing normative constraints of rationality. The negative claim holds that *it is not possible* to easily reduce a group's attitudes to its members' attitudes (See List and Pettit 2011, p. 4). Thus group agents are autonomous from and yet dependent upon their members. I briefly introduce their argumentation for the positive claim in this chapter. Yet, in particular their negative claim is important for a discussion of how group agents may get to believe a proposition in a justified way. Further, List and Pettit's discussion of this topic is of general relevance for our epistemic endeavour. Thus the negative claim should be given the necessary space. Significantly more space will be dedicated to this claim accordingly.

3.1. Agency and Rationality: The Positive Claim

List and Pettit foster a rather deflationary version of agency. In essence, even simple systems can, under certain circumstances, qualify as agents. According to Dennettian-like intentional stance theory,⁴⁶ to which List and Pettit are committed, if it makes sense to ascribe certain features significant for intentionality to an entity, this entity may rightly be

⁴⁵ Their version of joint intention is roughly comparable to Bratman's account of shared intentions. See List and Pettit (2011); Bratman (1999); and Pettit and Schweikard (2006).

⁴⁶ See Dennett (1987 and 1971); Pettit (2007).

referred to as an intentional system.⁴⁷ These features are representational and motivational attitudes and further include the capacity to process these attitudes and act in accordance with them. Representational attitudes “depict how things are in the environment” (List & Pettit 2011, p. 20) and hence they are belief-like states or judgements. Motivational attitudes “specify how [a system] requires things to be in the environment” (ibid), hence they are desire-like states or preferences.⁴⁸ The third feature regards a system’s capacity to not only process these attitudes, but moreover to act in accordance with them. In particular, this third feature points to a further criterion for sensibly ascribing agency to an intentional system, namely that minimal requirements of rationality need to be in place. List and Pettit categorize these requirements as belonging to three groups: attitude-to-fact, attitude-to-attitude and attitude-to-action requirements for rational agency. These requirements apply to both simple agential systems, like chess computers, robot vacuum cleaners or, as we shall see, groups, and more complex agential systems like human beings. However, in complex agential systems, like humans, these requirements are not sufficient for rationality. The demarcation criterion for simple and complex systems is their capacity to reason (in complex systems) or their lack thereof (in simple systems).

The first, attitude-to-fact, requirement refers to “standards of how to look out for evidence, how to respond to perceptual evidence in forming perceptual representations” (ibid. p. 24) and the like. In a nutshell, this requirement demands that attitudes are formed in accordance with ones environmental circumstances.⁴⁹ While this is probably all that can be demanded from simple systems, in complex systems attitude-to-fact rationality requires further restrictions that aim at an agents’ capacity to reason. While groups themselves cannot reason, they can get pretty close to reasoning, as we shall see below.

The second requirement for rational agency, attitude-to-action, rules out phenomena-like weakness of will or otherwise irrational actions. This requirement imposes the

⁴⁷ As long as certain ceiling and floor constraints are met (see List and Pettit 2011, p. 29).

⁴⁸ For a further discussion of the relation between beliefs and judgements on the one hand and desires and preferences on the other, see List and Pettit (2011, esp. pp. 25ff).

⁴⁹ This forms true rather than false beliefs, realizable, rather than unrealizable desires and intentions, and so on.

restriction of only performing actions in accordance with the attitudes held. It should warrant that only actions are favoured and performed that are “required – or at least permitted – by the agents’ representations and motivations and rule against those that aren’t” (ibid).

Finally, attitudes-to-attitudes rationality must be present or at least possible in intentional systems that qualify as agents. This requirement aims at two restrictions. First, it rules out instances in which means-end reasoning fails. Second, it rules out the possibility of holding inconsistent attitudes. Hence attitude-to-attitude requirements of rationality impose certain criteria of consistency onto an intentional system. This requirement is of particular import when discussing the rationality and autonomy of group agents.

By sensibly ascribing intentional states to a certain system and imposing rational restrictions on its behaviour, both allow us to refer to a system as an agent. While human individuals are highly complex systems capable of reasoning, groups, for example, do not by themselves exceed the above criteria and hence qualify as agents in this rather deflationary manner.⁵⁰ In addition, to the extent that meeting these rationality requirements demands active measures to ensure it, a group agent’s rationality will depend on its members.

Note that allowing for some groups to qualify as simple agents does not mean that their inner structure is simple. Quite the opposite is the case. This becomes clear when looking at the second claim that List and Pettit defend, that is to say, their negative claim. According to this, a group’s attitudes cannot in principle be reduced to the input attitudes of its members.

Just because it might be sensible to ascribe agency to some groups, this does not yet say that we can therefore state anything significant about them that we could not also formulate in terms of members’ attitudes, rationality and actions. In other words, merely granting that groups may qualify as agents does not yet allow us to state that such groups also qualify as autonomous with respect to their members’ attitudes. Accordingly, more

⁵⁰ Although there are certain specifics that clearly set groups apart from, for example, simple robots and where a group can resemble reasoning qua its members. The extent to which groups can “reason” should become clear in subsequent chapters.

needs to be said for supporting the idea that some groups not only qualify as rational agents, but that they also qualify as autonomous with respect to their members' attitudes and hence allow for claiming that NRR holds for groups. List and Pettit's negative claim entails a reason for the irreducibility of group attitudes with respect to member attitudes. I now turn to the negative claim.

3.2. A Theory of Attitude Aggregation: The Negative Claim

According to NRR, groups that qualify as rational agents in the deflationary sense mentioned above also qualify as autonomous with respect to their members' attitudes. Yet at the same time, a group's attitudes are comprised of those of its members. This autonomy therefore hinges on the transition of (individual) members' attitudes to the (corporate) group level.⁵¹ However, this transition is not only necessary for a group to hold any kind of attitude, but also fails to warrant reducibility to members' attitudes. In this sense, the autonomy claim of group agents with respect to their members depends on making a case in favour of the idea that a group's attitudes are not easily reducible to those of its members. The making of such a case constitutes the negative claim.

Before discussing how member attitudes determine group attitudes, as well as how this specific relationship nonetheless warrants the autonomy of groups, three further remarks are in order. First, the feature that warrants this irreducibility is not only the specific relationship between member attitudes and group attitudes, but also a group's "organisational structure" (see List and Pettit 2011, esp. pp. 59ff).

The organisational structure of a group covers various aspects that can hardly be captured in a somewhat technical discussion derived from social choice theory, which gives rise to List and Pettit's irreducibility claim. A group's organisational structure not only covers aspects like whether a group is democratically organised or hierarchical, but also includes, for example, the mechanisms and practices that relate to how potentially democratic or hierarchical structures are implemented.

⁵¹ Recall Figure 1 in Chapter 1. The autonomy of group agents hinges on the upper arrow depicted there (see p. 32).

What we can state for now is that, second, in practice the relationship between member attitudes and a resulting group attitude comes in very different forms. However, the irreducibility claim of group attitudes with respect to member attitudes is based on a formal result in judgement aggregation theory and not on practical grounds. Therefore, while I approach this issue by way of a practical example, these examples should simply make more intelligible the formal element. It is important to keep in mind that the extent to which such formal results do or should map onto or diverge from reality remains an open question.

Third, some technical terms need to be made clear from the outset. For one, attitudes are construed (as List and Pettit suggest) as adopting a “positive” or “negative” stance towards a proposition. Such a proposition might concern either preferences or beliefs. In each case, these attitudes comprise “Yes” or “No” judgements and are thus of an on-off binary kind (see *ibid*, p. 25ff)⁵². Furthermore, in what follows an aggregation of judgements is supposed to mean any kind of meshing of member judgements such that a single corporate judgement is the result, unless further specified. In general, it is that which makes out of individual stuff, corporate collective stuff. In accordance with List and Pettit’s discussions, propositions towards which member judgements are aggregated are referred to as the group’s *agenda*, the different opinions of the members regarding propositions on the agenda in the course of an aggregation of judgements taken together is henceforth called the member *profile* (cf. *ibid*, esp. p. 48f). With these brief remarks to hand, it is time to discuss List and Pettit’s negative claim in greater detail and with it the theory of judgement aggregation.

Aggregation Functions and a Group’s Organisational Structure

Every potential group agent (or rather, its members) implements certain procedures for adding up members’ attitudes in order to arrive at a group attitude; this is essential for acting upon these attitudes and thus for group agency. In standard social choice theory, such procedures are referred to as aggregations of judgements. In technical terms, we are able to determine certain aggregation functions, for example majority voting, expert panel decisions, unanimity votes or the dictatorships of one or several members. Such functions

⁵² A similar discussion can be found in Pettit (1993).

can, in practice, be implemented explicitly, such as in groups where certain by-laws determine how a group comes to hold a certain attitude (think of parliaments, for example). Alternatively, such mechanisms can also be implemented in an inexplicit way. Inexplicit mechanisms often involve deliberations among members and may further allow for adjustments when necessary due to feedback loops between resulting group attitudes and the member profile. In other words, in the aggregation of judgements members may implement the possibility to check on the resulting group attitudes in order to identify certain potential breaches of rationality on a group level. The advantage of such “inexplicit” modes for aggregating judgements is that a group is more flexible when determining its attitudes and can thus adjust its aggregation procedures according to the member profile.⁵³

In order to become clear on the role that aggregation processes play and to gain a better understanding of different mechanisms and what they each involve, an example will be of help. Thus the robbers that, in the first chapter, were queuing in the coffee shop, not to buy a hot beverage but in order to rob the shop, are discussed here once again.

Given that the collection of robbers forms a group agent, the members will have decided upon various issues. Among these is, for example, whether or not they want to rob coffee shop C at a certain time. Say they do rob C. How will the decision upon which the group acts have come about? There are various scenarios that we might imagine to have taken place. It could be that the group has one boss who is in the position to always decide what the group will do. This is what List and Pettit refer to as the “dictatorship” of one member (see, for example, *ibid* p. 36 or List 2005). Or maybe there are two bosses that together determine the group’s “mind”. In such a case, we would be confronted with a dictatorship of a certain subgroup within our group of robbers. Alternatively, we might imagine that the group is organized in a rather democratic manner. In this case, the members will have discussed the advantages and disadvantages of robbing the coffee shop, and then have voted on the matter, or maybe they are not so much interested in a discussion and simply vote over the issue in silence by raising their hands or perhaps even

⁵³ However, this flexibility can be a disadvantage as well, see List and Pettit (2011, esp. pp. 57f and pp. 104ff).

casting ballots. Whatever it was that leads to the group's attitude (in this case, to wanting or intending to rob C), according to aggregation accounts what all these alternatives share is the fact that there exist a number of (or in extreme cases only one) opinions regarding whether or not to rob C that the members (or the boss) hold, and these attitudes are somehow summed up in the group's attitude and thus determine what it does. In short, there exists some input set of judgements as to whether or not to rob C and these judgements are somehow put together in such a manner that they result in an output, namely the group's position regarding whether or not to rob C. In addition, the formal procedure by which the input attitudes are added up, that is, aggregated to the group's attitude, is what List and Pettit call an aggregation function. If our robbers had chosen a majority vote or dictatorship as their general aggregation function for forming group attitudes, and if the aggregation were performed without any further deliberation on the group judgement, their aggregation procedures would be functionally explicit. But in the case that our robbers discuss the matter first, that is, deliberate together over whether or not to rob C, we are confronted with a rather inexplicit aggregation procedure in which some "heuristic, from proposition to proposition" (List and Pettit 2011, p. 61) is involved.

Aggregation Functions and the Discursive Dilemma

Apart from functionally explicit and inexplicit aggregation procedures, some more general remarks regarding aggregation functions are in order. A simple figure, inspired by List and Pettit, should clarify the general idea (Fig. 2).

In general terms, more or less anything that turns a set of member attitudes into a group attitude could serve as an aggregation function. However, we assume that our robbers are rational (yet probably immoral) agents, or at least that they should be rational, and so only certain aggregation functions are sensible.⁵⁴ As such, tossing a coin will probably not suffice. In fact, there seem to be some minimal requirements that appear to be commonsensical and should ensure rationality on a group level. List and Pettit present

⁵⁴ Remember that the methodology that List and Pettit use is one in which a group only qualifies as an agent if it meets certain criteria of rationality. Hence, if our robbers fail to meet these requirements, the group will not qualify as an agent.

four conditions that an aggregation function has to fulfil in order to be “plausible” (cf. List and Pettit 2011, p. 49).⁵⁵

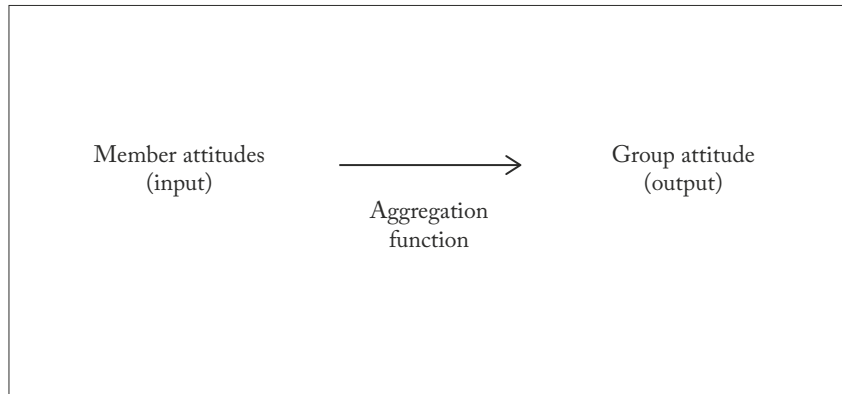


Figure 2 List and Pettit’s Model of Aggregation Functions (cf. List and Pettit 2006 and 2011)

- (1) **Universal domain**, which essentially holds that an aggregation function must be able to deal with any set of input that is “consistent and complete”⁵⁶ (ibid.).
- (2) **Collective rationality** is required, that is, the group’s (i.e. output) attitudes at least have to be consistent and complete.
- (3) **Anonymity** requires, roughly, that each individual’s attitude be given the same weight.
- (4) **Systematicity** involves two aspects. First, it requires that the group attitude towards a proposition derive from the member attitudes towards that same proposition only (independence condition). Second, it requires that “the same pattern of dependence should hold for all propositions” (ibid, p. 54).

Perhaps the most obvious candidate for an aggregation function that meets these four criteria is simple majority voting. If all the above conditions could be met by a simple majority aggregation procedure then List and Pettit could refer to it as a “redundant”

⁵⁵ In some of List and Pettit’s papers, as well as other contributions to the field of judgement aggregations, only three instead of four criteria are introduced. Best known, perhaps, is Arrow’s (Arrow 1963) introduction of these criteria (or a version of them). I stick to the most recent set, that presented in List and Pettit (2011).

⁵⁶ A set of judgements is consistent if the proportions held are co-realizable. A set of judgements is complete if the necessary implications are also determined. Cf. List and Pettit (2011, esp. p. 49)

relation. This is so because the resulting group attitude could be directly reduced to the members' attitudes. In such a case, if we knew the individuals' attitudes then we would know what the group's attitudes were. If it were so easy then we could probably place to one side any consideration or discussion of the irreducibility of groups.⁵⁷ But this simple aggregation is blocked and so a case of irreducibility of group attitudes to member attitudes ensues. Perhaps the most important argument as to why it is blocked derives from discursive dilemmas. Table 1 shows one version of the dilemma.⁵⁸

An example is once again of help. Take our robbers. Say that our group of robbers counts three members and they are to judge upon the following:

P: Is it easy enough to rob C?

Q: Does the financial pay off of robbing C suffice?

$P \wedge Q$: Should we rob C? (Given that P and Q are necessary and jointly sufficient to settle this matter.)

Let us further assume that what follows constitutes the robbers' set of attitudes towards these propositions (which are in themselves consistent).

	P?	Q?	$P \wedge Q?$
Member 1	Yes	No	No
Member 2	No	Yes	No
Member 3	Yes	Yes	Yes
Aggregation/Group	Yes	Yes	Yes / No

Table 1 Discursive Dilemma

Discursive dilemmas show one way in which a group's judgement may end up contradicting the majority's judgements, such that reducibility is blocked. In essence, this is so because it is unclear as to whether the group attitude is "Yes" or "No" in the first

⁵⁷ Pace Sylvan (2012).

⁵⁸ I will not explain in detail what goes on in discursive dilemmas. They have been introduced often enough. Rather, I will elaborate upon them only to the extent needed for the overall project of this thesis. For an overview of discursive dilemmas and group agency see, for example Pettit (2010) and List and Pettit (2002).

place – it depends on whether we count the premises (that is, P, Q) or the conclusion (that is, $P \wedge Q$) (see Table 1). No matter whether the group adopts that yes $P \wedge Q$ or not, in any case the group members' opinions will not be sufficiently accounted for and hence the aggregation function will have problems in accommodating this input set of judgements, which violates (1). Above all, the group holds both yes and no concerning $P \wedge Q$, and this clearly violates (2).

How, then, could such a scenario be avoided without violating any of the above conditions? We could in fact construct various scenarios and each of them would result in a violation of one or the other of the above conditions. This is essentially what List and Pettit's impossibility theorem (cf. List and Pettit 2011, p. 50 and List and Pettit 2002) holds. Moreover, it is what List and Pettit echo in their NRR account, since this discrepancy provides them with reason to assume that a simple reduction from a group's attitude to the members' attitudes does not, in principle, work.⁵⁹ This being so, how are we to deal with the discursive dilemma of our robbers? List and Pettit suggest that in many contexts it would be reasonable to relax (though not fundamentally violate) (4). In this particular example, there exist two ways out for our robbers. One is to go for a premise-based procedure in which P and Q are added up and deduction decides the rest. In this case, the robbers would rob C and (4) is relaxed.⁶⁰ The other option would be to only vote on the conclusion. This solution, however, comes with further problems, in particular with respect to the consistency of the group's set of attitudes in the long run. Moreover, conclusion-based procedures violate the completeness criterion in (1), since judgements on the premises would not be considered, despite being part of the group's agenda. Hence List and Pettit suggest that premise-based procedures or, more generally, sequential priority procedures are to be favoured over alternative options.⁶¹

⁵⁹ Another reason as to why the reduction does not work according to List and Pettit is that a group's organisational structure is a complex thing, determining which aggregation functions we choose, which way out of a discursive dilemma we favour, and how we cast our votes and various other practices that may play into the aggregation of judgements (see List and Pettit 2011).

⁶⁰ Note that "relaxing", for example, (4) does not necessarily imply that we also lose group rationality. In fact, there can be reasons to relax each of these criteria depending on whether or not it is a good idea in a certain context. Again, I will not deal with this in greater detail here, the mere claim must suffice (cf. List and Pettit 2011).

⁶¹ Sequential priority procedures are characterised by the inclusion of a ranking of the propositions on a group's agenda. Such a ranking can consist in viewing some propositions as premises and others as conclusions, or otherwise in prioritising some propositions over others. By means of prioritisation, potential

While it is clear that in such procedures member attitudes still determine the group's attitude, it also becomes evident that these two come apart in significant ways. Leaving out the technicalities of premise-based procedures, it is important to note why systematicity is violated. The reason is that the conclusion (in the example above to yes, rob C) is not derived from the members' attitudes regarding whether or not to rob C (as systematicity would demand it) but from the premises for this judgement. This is one particularly important sense in which a group's attitude may not directly map onto those of its members. The relationship between member judgements and those of the group is, List and Pettit conclude, that of set-wise, or holistic, supervenience (List and Pettit 2006). A little more needs to be said about set-wise supervenience since this is what motivates the negative claim, namely that a group's attitude, if it is supposed to be as rational as possible, are not directly reducible to the attitudes of the members.

Holistic Supervenience

Supervenience is a common term within discussions in the philosophy of mind. Albeit opaque at times what can be said about supervenience is that it diverges from causality in several important respects (cf. Stalnaker 1996). For example, in classical accounts of causality, it is often assumed that causing occurs temporarily prior to that which is caused. This need not be the case in supervenience relations. Moreover, it is not a necessary feature of supervenience relations that one thing makes it the case that another thing occurs. In this respect, supervenience denotes a different kind of connection between one thing and another. The primary feature of supervenience is to state two features of a relationship between two entities or events. For reasons of simplicity, let us assume that there is an A set and a B set where the B set supervenes upon the A set. The first feature of supervenience relations is that changes in the B set can only occur given changes in the A set. However, this does not say anything about a temporal relationship between the A set and the B set, that is, the changes in the A set would have to occur prior to changes in the B set. They might just as well occur simultaneously. Furthermore, supervenience is characterised by multiple realisability. Variations in the A set need not necessarily involve changes in the B set and so there may exist various A sets that make for the same B set.

inconsistencies, as well as other unwelcomed rationality breaches on a group level, can be avoided. Premise-based procedures are but one example of sequential priority procedures. See, for example, List (2004a).

List and Pettit's example is of a cluster of dots on a piece of paper which, when looking at them, cause us to see a particular shape. There can be variations with respect to the dots without the shape actually changing. Some of the dots could be at slightly different points on the paper, or we may switch certain dots around. Some could be bigger in one case than in another. Still, the shape could remain the same (see *ibid*, esp. pp. 65ff).

With respect to the relationship between member attitudes and a group's attitudes, the supervenience claim equally holds that in order for changes to occur in the group set, there need be changes in the member set. However, there might be changes in the member set that do not alter the group set. In particular, the supervenience relation that List and Pettit assume for the relationship between member attitudes and group attitudes is further characterised by being holistic in nature (see List and Pettit 2011, pp. 65ff). Holistic supervenience means that group attitudes not only supervene upon each and every member attitude in isolation, but that the group attitudes do so with respect to the member profile as a whole; in other words, the "set of group attitudes across propositions is determined by the individual sets of attitudes across these propositions" (*ibid* p. 69). This can be illustrated by the premise-based or, more generally, the sequential priority procedure approach to aggregating judgements. In Table 1, we can see that a premise-based procedure would yield the result that, yes, C will be robbed. This follows from the majority position on the premises P and Q. However, if we aggregate the members' attitudes towards the conclusion $P \wedge Q$ then the group would hold that, no, C will not be robbed. As illustrated, both the conclusion and the premise-based approach would neglect certain attitudes of the members. In the case of the conclusion-based aggregation, the majority position towards the premises would have to be ignored. In the case of the premise-based aggregation, the majority position towards the conclusion ($P \wedge Q$) would have to be ignored. In the premise-based aggregation, however, it is also the case that the conclusion is not a direct result of the members' attitudes towards that particular proposition, rather it simply follows from the premises. In this sense, the group's attitude comes apart from the members' attitudes towards the proposition $P \wedge Q$. This is so because the whole profile is understood to determine the group position, in this case giving priority to the premises rather than to the conclusion. Hence the conclusion that yes, we will rob C, depends holistically on the members' profile towards the premises.

Holistic supervenience, which is dependent on certain aggregation functions, is thus key to claiming NRR about groups. The NRR account assumes that while member attitudes do not simply or in any direct way determine the group attitude, the group attitude supervenes holistically upon the attitudes of the members.

Let us briefly recap the essential features of NRR about groups. First, groups can, in principle, perform as rational agents, to the extent they can hold desire-like, belief-like and intention-like states, and they can process these states in a minimally rational way. This is List and Pettit's positive claim.

Second, aggregation functions and the procedures for implementing them are a necessary means for groups to hold attitudes. An aggregation function denotes (the formal side of) various means by which the individual judgements of the members are meshed or merged so as to result in one corporate judgement. However, no aggregation function allows, in principle, for rational group attitudes without breaching one or the other requirement that should secure rationality on the group level. This can be illustrated by discursive dilemmas and other paradoxes concerning the aggregation of judgements or preferences, which leads to various impossibility results in general, and List and Pettit's impossibility results regarding aggregation functions in particular.⁶² Yet there are sensible ways in which to deal with these results and thereby to sustain rationality.

Third, motivated by the ways in which to cope with the impossibility theorem about judgement aggregation in order to secure group rationality, group attitudes are best conceived of as supervening holistically upon members' attitudes, rather than in a straightforward proposition-wise manner. Consequently, in addition to or as part of the general complexity of a group's organisational structure, and in light of the solutions regarding how to overcome the impossibility theorem of judgement aggregation, a simple reduction of a group's attitude to members' attitudes is not possible. This is the negative claim in favour of NRR, as introduced at the beginning of this chapter. Accordingly, groups are best conceived of as autonomous with respect to their members' attitudes.

⁶² See for example, Arrow (1963); List and Pettit (2002); and (2004).

Having presented a detailed introduction to List and Pettit's account of group agents in general, the epistemic implications of this theory should be considered. This is what the following chapter aims to do.

Chapter 4

Meeting the Knowledge Challenge. List and Pettit's Proposal

So what about justified belief or knowledge in List and Pettit's groups? To begin with, some preliminary comments are in order. Cases in which members know something do not automatically count as group knowledge, even if the majority of a group's members know that, say, *p*. This should be clear from the Introduction and the first part of this thesis. Alternatively, what we are looking for in List and Pettit's account is an attitude, and more precisely a belief, that is held by the group and that is justifiably held in that it has been reliably produced. As seen in the last part of this chapter, the way by which a group agent can come to hold attitudes involves some kind of aggregation function. It is thus natural to assume that the transformation of member beliefs⁶³ to a corporately held group belief by means of an aggregation of judgements is precisely what interests us here. We also have to assume, of course, that certain propositions are aggregated, the contents of which are thus either true or false. Henceforth, I call beliefs regarding facts simply belief states, as opposed to states like desires or intentions, and also as opposed to means-ends beliefs and other belief-like states that are not usually the primary targets of epistemological discussions. Given that aggregation functions provide group attitudes, belief states that qualify as justified or even as knowledge also need to be aggregated somehow and depend on a group's member inputs. Hence the focus of our discussion lies in the transition of belief (or, rather judgements) from the member level, to the group level. Hence, we find ourselves in the midst of what has been suggested as the central issue matter for collective epistemology in chapter 2.

Considering List and Pettit's account, what we need to find is an aggregation function, or a set of aggregation functions, that provide(s) us with the possibility of the outcome belief being true and indeed being so reliably. List and Pettit, and in particular List, have a story to tell within their NRR account about how the "knowledge challenge"

⁶³ Or rather judgements. See below.

(List 2005) for groups can be met.⁶⁴ Accordingly, in this particular chapter I examine their discussion about how a group may come to hold beliefs that aim at truths.

Since the following discussion is fairly technical, in order to make plausible the idea of group knowledge, let us first return to another example that involves our three robbers. This will serve as a proper introduction to the discussion that follows. Let us suppose that our robbers must work out a certain fact. Let us say that the group has to make up its mind as to whether or not there is a surveillance camera installed in the coffee shop. Let us call the proposition that *there is a surveillance camera in the coffee shop*, *p*. Let us further assume that *p* is true and that the group comes to believe that *p*.

Regarding the case of our robbers, we must ask whether and to what extent various aggregation functions that our robbers may use determine the resultant group belief and if yes, to what extent. Furthermore, are there more preferable routes for aggregating beliefs than others? Also, we will have to ask ourselves how a group belief can be reliable in light of its dependence on the member beliefs.

As already mentioned, a discussion of a group's potential performance as an epistemically rational agent includes some rather technical deliberations. These draw on various fields and in particular social choice theory, Bayesianism and probabilistic truth-tracking accounts. I discuss elements of the vast and related field of literature only to the extent needed for the aim herein of tracking down the possibility of group knowledge. First, however, let us briefly recap what has been learnt so far.

Impossibly results from judgement aggregation show that there is no aggregation function that could, in principle, allow for meeting the four criteria introduced in Chapter 3. This observation results in a brief discussion of the suggestions that List and Pettit make regarding how this impossibility result might be compensated for in order, nevertheless, to ensure a group's rationality (see esp. List and Pettit 2011, pp. 81ff). We have seen that different aggregation functions, such as dictatorship, majority vote, conclusion-based procedures or premise-based procedures (as a special case of selective-

⁶⁴ Note that while up to now the discussion has most often pertained directly to List and Pettit (2011), List has published on this issue rather extensively. Hence in the following discussions I draw upon various papers, especially List's. Thereby, I suppose that List's contributions to these topics are coherent with List and Pettit's NRR account in their 2011 book.

priority procedures) each violate different conditions, but that some aggregation procedures nevertheless perform well.

While simple rationality requirements are of course important, the epistemic justification of corporate beliefs tells a different story. How can a group perform in such a way that it forms a corporate belief that reliably tracks truth? The simple yet challenging answer is, it depends, and, “there may not be a ‘one size fits all’ organisational design that is best for all group agents and all epistemic tasks” (List and Pettit 2011, p. 102).

However, we are able to make a general case in favour of the possibility of group knowledge. In judgement aggregation theory there exist possibility results that allow us to make a convincing case that groups can hold beliefs that successfully track truths. The most persuasive result is Condorcet’s Jury Theorem (CJT). It holds that under certain formal conditions, and given a choice between two options (say p and $\neg p$) where one is “objectively” true or correct, the probability that the majority will pick the correct one of the two options tends towards 1, as the group size increases. For this to hold, it suffices that the member judgements exhibit some, albeit very slight, bias towards the truth.⁶⁵ In addition, given that the members’ judgements are equally reliable, which is indeed part of one of these conditions, the majority judgement will be more reliable than any of the members’ judgements. Accordingly, groups can, under certain circumstances, perform well in “tracking truths” (cf. Nozick 1981), and indeed do so even better than any of their members.

The formal conditions of CJT referred to above are *independence* and *competence*. To better understand these two conditions, and the epistemological context discussed here in general, some basic vocabulary from probabilistic theory, and indeed List and Pettit’s adoption thereof, is of use.

⁶⁵ This is one implication of the much wider phenomenon of the “law of large numbers”. See, for example, Hacking (2001, esp. pp. 205ff). Also note that Condorcet’s Jury Theorem goes beyond “correctness” as an aim. See Estlund (1994).

4.1. Truth-Tracking and Probabilities of Judgement

List and Pettit, and in particular List (2005), suggest a Nozickian framework of truth-tracking in order to formulate their possibility result with respect to CJT. According to Nozick, there are four conditions that are necessary and jointly sufficient for knowledge. Those are (Nozick 1981, pp. 172ff):

- (1) p is true
- (2) S believes that p
- (3) If p weren't true, S wouldn't believe that p
- (4) If p were true, S would believe that p ⁶⁶

Various refinements and objections have been presented in reaction to Nozick's four conditions of knowledge.⁶⁷

According to Nozick, (3) denotes the negative tracking reliability of a subject, S, and (4) constitutes S's positive tracking reliability. Since (1) and (2) are rather self-explanatory and simply denote an externalist framework, conditions (3) and (4) are the relevant ones. The idea is that S possesses knowledge if in close possible worlds both (3) and (4) hold.⁶⁸ Condition (3) is of particular import here. According to this condition, an agent knows that p only if she would (in close possible worlds) not believe that p , were p false. This constraint should exclude examples like that of the following. Imagine one of our robbers, Roberta. Recall the example above in which our robbers were to judge on whether p , namely *there is a surveillance camera in the coffee shop C*. Roberta sees an object that looks like a camera hanging from the ceiling and concludes that p and indeed p is the case, hence (1) and (2) are given. However, imagine an altered version of Roberta's case. In this scenario, Roberta believes that p , which is true, however, the camera she sees, and on the basis of which she forms her belief is fake. The belief is true nonetheless, because,

⁶⁶ Note that this is only the shortened version of his conditions for knowledge. The more detailed version includes the method by which S holds her believe that p in order to meet various counterexamples that would occur if sticking to this general version. See Nozick (1981 p.179ff). See also Chapter 6 below for a discussion.

⁶⁷ See, for example, Kripke (2011), or Goldman (1983).

⁶⁸ How "close" these possible worlds are supposed to be, especially regarding condition (3), is one of the challenges for this account. To provide a clear account of which possible worlds should be considered as relevant for (3) has been shown to be difficult, if at all, determinable. See Goldman (1983) See also Chapter 5 below.

unbeknown to her, a real surveillance camera is hidden in the droplight. Accordingly, p is true, and Roberta truly believes p . However, we would not hold that she knows that p . Indeed, condition (3) conveniently allows us to exclude this case, since in close possible worlds in which there is no hidden camera, and so where p is not true, Roberta would still believe that there is a surveillance camera in the coffee shop. This violates condition (3) and hence Roberta does not know that p .

List's adoption of this framework (List 2005, see also List and Pettit 2011, esp. pp. 81ff) includes a probabilistic "twist". Rather than stating that, given (1) and (2), S knows if (3) and (4) are met in all or most close possible worlds, the idea is to assign a certain credibility to the truth-tracking ability of an agent. Accordingly, List reformulates Nozick's conditions for positive and negative truth-tracking (List and Pettit 2011, pp. 81ff, see also (List 2005). What we are looking for to ensure truth-tracking is the following:

- (I) Positive tracking reliability: The conditional probability that the agent judges that p , given that p is true.
- (II) Negative tracking reliability: The conditional probability that the agent does not judge that p , given that p is false.

(I) and (II) denote the probabilistic version of the subject's positive and negative truth-tracking reliability (List 2005, pp. 31ff, List and Pettit 2011, pp. 81ff), where (II) corresponds with Nozick's condition (3), and (I) with his condition (4).

When formalised (I) can be written as $\Pr(J|T)$ and (II) $\Pr(\neg J|\neg T)$, where J denotes the agent's judgement that p and T denotes the fact that p is true. The higher the probability, the better a subject is at truth-tracking. However, to track the truth rather than not, the probability must be greater than 0.5.

In order to determine $\Pr(J|T)$, we can insert it into the standard version of Bayesian Theorem:

$$\Pr(J|T) = \frac{\Pr(J) \Pr(T|J)}{\Pr(T)}$$

where

$$\Pr(T) = \Pr(J) \Pr(T|J) + \Pr(\neg J) \Pr(T|\neg J)$$

This substitution presupposes that $\Pr(J) = 1 - \Pr(\neg J)$.⁶⁹ We see that in order to determine $\Pr(J|T)$ we need to have information on $\Pr(T|J)$ and $\Pr(\neg T|\neg J)$ ⁷⁰ in addition to the probability $\Pr(J)$. $\Pr(T|J)$ and $\Pr(\neg T|\neg J)$ are precisely the reverse of truth-tracking reliability and constitute what List calls “indicating reliability” (see List and Pettit 2011, esp. pp. 81ff). When tracking reliability amounts to the conditional probability of the agent’s judgement that p given p is true and the conditional probability of the agent’s judgement that not p , given p is false, indication reliability tracks the reverse. It therefore tracks the conditional probability of the proposition being true given that the agent judges that p is true, as well as the conditional probability of the proposition being false given that the agent judges it to be false. The table in Table 2 summarises the conditions and related probabilities of both indicating and truth-tracking reliability.

Condition of truth-tracking	Condition of truth indication
Positive: If ‘ p ’ were true, the agent would judge that p .	Positive: If the agent were to judge that p , ‘ p ’ would be true
Negative: If p were false, the agent would not judge that p	Negative: If the agent were not to judge that p , p would be false
Corresponding probabilities	Corresponding probabilities
Positive: The conditional probability that the agent judges that p , given that p is true; hence $\Pr(J T)$	Positive: The conditional probability that p is true, given that the agent judges that p ; $\Pr(T J)$
Negative: The conditional probability that the agent does not judge that p , given that p is false; hence $\Pr(\neg J \neg T)$	Negative: The conditional probability that p is false, given that the agent does not judge that p ; hence $\Pr(\neg T \neg J)$

Table 2. Positive and Negative Truth-Tracking and Indicating Reliabilities

⁶⁹ That is to say, since $\Pr(J)$ denotes the probability that an agent judges that “ p ”, and $\Pr(\neg J)$ denotes the probability with which she judges that “ $\neg p$ ”, it presupposes that the members do not suspend judgements but rather, in each and every case, either judge that “ p ” or judge that “ $\neg p$ ”.

⁷⁰ We need information on either $\Pr(\neg J|\neg T)$ or $\Pr(J|T)$, because $\Pr(J|T) = 1 - \Pr(\neg J|\neg T)$.

According to List and Pettit, the following can be inserted in the above probabilities: $\Pr(J|T)$ is the agent's positive truth-tracking reliability; $\Pr(\neg J|\neg T)$ is the agent's negative truth-tracking reliability; $\Pr(T|J)$ is the agent's positive-indicating reliability; $\Pr(\neg T|\neg J)$ is the agent's negative-indicating reliability; $\Pr(J)$ and $\Pr(T)$ are the unconditional probabilities of judging that p in case of $\Pr(J)$ and that p is true, in case of $\Pr(T)$.

Just as we can determine an agent's truth indication $\Pr(J|T)$ with respect to p by inserting it into the above theorem, so we can determine an agent's indicating reliability by inserting it into the same formula. Accordingly:

$$\Pr(T|J) = \frac{\Pr(T) \Pr(J|T)}{\Pr(J)}$$

where

$$\Pr(J) = \Pr(T) \Pr(J|T) + \Pr(\neg T) \Pr(J|\neg T)$$

We see that $\Pr(T|J)$ depends on probability $\Pr(T)$, hence the probability with which p is true; on $\Pr(J|T)$, hence an agent's positive truth-tracking reliability; and $\Pr(\neg J|\neg T)$, hence an agent's negative truth-tracking reliability.

Note that while $\Pr(J|T)$, which is an agent's positive tracking reliability, and $\Pr(\neg J|\neg T)$, which is an agent's negative tracking reliability, can coincide, this need not be the case. In fact, they often do not do so.⁷¹ An example will help to clarify this. Let us again refer to Roberta. Roberta may have a high positive tracking reliability of, say, 0.98 with respect to p . This means that her chance of judging p is very high and, correspondingly, that her chance of judging that $\neg p$ although p is the case – that is, $\Pr(\neg J|T)$ – is very low, namely $1-0.98=0.02$. However, this does not yet indicate that her negative tracking reliability is equally high. She might believe that nearly everything hanging from the ceiling is a surveillance camera or she might be slightly paranoid and believe that in every hanging light a camera is hidden. Hence her negative tracking reliability may be very low, since she judges that p in practically all circumstances, even in the absence of a camera. For example, she might falsely attest that a surveillance camera is present, given that it were not, that is $\Pr(J|\neg T)$ with a probability of 0.86. Accordingly,

⁷¹ The same applies to an agent's positive and negative indicating reliability.

her negative tracking reliability, $\Pr(\neg J|\neg T)$, would only be $1-0.86$, hence, 0.14 . Framed using a different vocabulary, she is good at avoiding false negatives but very bad in avoiding false positives.

With these probabilistic tools to hand we can now return to List and Pettit's adoption of CJT and, in particular, its *independence* and *credibility* constraints.

Usually, the independence of judgements is warranted if the following applies: $\Pr(J^1|J^2)=\Pr(J^1)$ ⁷², where, in our case, $\Pr(J^1)$ denotes the probability of one member's judgement that p and $\Pr(J^2)$ denotes another member's judgement that p . In the example of Roberta, independence assumes that the probability attached to her judgement that p does not depend on some other member's judgement regarding p .

The second condition regards *competence*. This condition consists of two parts. First, it is assumed that the credibility of the members must be better than arbitrary. In the vocabulary introduced above, this amounts to the claim that the members must have a greater positive (and negative) tracking reliability than average, that is to say, greater than 0.5 .⁷³ The second part of CJT holds that all members must have the same tracking reliability, both positive and negative.

Given that these two conditions apply, the following can be shown. The group's probability of correctly judging on whether p increases and indeed tends towards 1 – and hence certainty – with an increase in the number of members who judge on p and given an aggregation of judgements by means of majority rule. At least under these favourable conditions, in which the law of large numbers works for the group's cause, we can make a strong case in favour of a group's capacity to form true beliefs. If CJT applies then the bigger the group, the more a group's judgement tends towards certainty. By making this effect work for a group we are able to obtain “epistemic gains from democratization” (List and Pettit 2011 p. 87ff).

Of course, this is a mere possibility result. In actual practice, not all of the requirements necessary to meet CJT will be met. More than often none of these criteria will be in place. Still, there are various ways in which a group can perform epistemically

⁷² See for example Estlund (1994, p. 138).

⁷³ However, we will see below that there exist scenarios in which it is not the case that both, negative and positive tracking reliability must be above the threshold of 0.5 .

successful even if the above criteria are not met. A closer look at some of these variations and the extent to which there still exists the possibility of epistemic gains from judgement aggregation is now in order.

4.2. Epistemically Promising Strategies of Judgment Aggregation in Groups

Indeed, there are various ways in which to criticise or spell out various aspects regarding CJT. Most obviously, deliberation among group members and judging based on a shared set of evidence could potentially threaten the independence constraint. However, various relaxations of the independence condition affect the overall effects of CJT to differing degrees, but each of these adaptations still allows for a weaker claim regarding a group's performance in truth-tracking, where the effect of a majority vote is weakened but still significant.⁷⁴ In other words, even under various conditions of relaxing independence there may still be epistemic gains from majority voting.

Equally, Condorcet's *competence* condition has been questioned in various ways. In usual cases of forming a corporate judgement based on the judgement of various members, the tracking (and indicating) reliability of the members varies, in many cases even significantly so, thus violating the second part of the competence condition, as well as, perhaps, the first. The result is that the subsequent group judgement will not necessarily be more reliable than those of each and every single member. Furthermore, the first part of the condition can be questioned. In fact, given that the second part does not hold in ordinary cases, we can very well imagine that some members in a group will be worse than average at judging a certain proposition, both regarding their positive and negative tracking (or indicating) reliability.

Note that, in fact, it is not necessary that both negative and positive tracking reliability (or indicating reliability) need be above the threshold of 0.5 for an agent to be a reliable truth tracker or indicator. For example, given that a member has a positive truth-tracking reliability of 0.4 – hence, below “arbitrary” – but a high negative truth-tracking reliability of, say, 0.8, it might still be the case that the agent is reliable in indicating the truth of p

⁷⁴ See, in particular, Dietrich and List (2004); Estlund (1994).

given her judgement. Whether or not this is so will also depend on the unconditional probability of p 's truth, $\Pr(T)$. In this case, if we assume that p is unlikely to be true, say, $\Pr(T) = 0.3$, her indicating reliability will still be above 0.5, namely 0.8.⁷⁵

Although this and related issues are widely discussed topics in judgement aggregation theory, I restrict myself to discussing those cases that List and Pettit mention and which indeed appear to be especially relevant to the aims of this paper. In the next chapter, I will introduce a special case in which negative and positive tracking reliability vary. This discussion will further motivate a process reliabilist framework instead of List and Pettit's truth-tracking account.

In relation to the competence condition, an interesting phenomenon explains how a member's tracking reliability (positive and negative) may fall below the required threshold of 0.5; that is to say, it may be worse than arbitrary (see List and Pettit 2011, esp. pp. 92ff). As seen in the brief discussion of discursive dilemmas, sometimes a group's agenda consists of various conjunctive propositions. Let us assume that there are n propositions, $p_1 - p_n$, and one conclusion, c . Let us further assume that the conclusion holds true only if either all the propositions $p_1 - p_n$ hold true (hence, the propositions logically hang together by means of a conjunction) or at least one of the propositions needs to hold true (hence the propositions are logically connected by means of disjunction). Given such a situation, it can be shown that a member's tracking reliability can easily fall below 0.5. This is so because even if we assume that a member has the same reliability towards each of the premises of, say, 0.6, her probability of tracking the truths of c will significantly drop since it is 0.6^n , which falls below 0.6 and will drop below 0.5 if $n \geq 2$. This observation alone should suffice to make yet another case in favour of premise-based aggregation procedures, since in this scenario the members only have to judge the premises. Formal results support this claim.⁷⁶ Consequently, this supports the idea that there can be "gains from decomposition" (List and Pettit 2011, p. 92).⁷⁷

⁷⁵ 0.795 to be exact.

⁷⁶ See List and Pettit (2011), pp. 94f

⁷⁷ Note, however, that Cariani (2012) in particular doubts whether it makes sense to derive the probability on a conclusion by means of exponentiating the tracking reliability of the member judgements on the premises. For his counter example, see Cariani (2012, p. 266f)

As mentioned above, usually members of a group do not have the same tracking reliability when compared to the tracking reliability of the other members and with respect to the various propositions on an agenda. The latter scenario is indeed necessarily so in the case of logically connected propositions like those presented in the example above. In order to provide a workable account of group knowledge, and in particular of the epistemic justification of group beliefs, these scenarios should be conceived of as the rule rather than the exception. The question is thus how a group can, ideally, perform given that the competence condition of CJT is violated due to a variation in the members' beliefs? In order to clarify this concern, an extreme and idealised case will be of assistance here. Imagine a group of 100 people who (in our idealised case) are scientists working together on a paper. They seek to figure out whether their hypothesis h is true or not. To determine this, they first have to decide on 100 premises. Each of the members is an expert regarding one and only one of these propositions. Imagine, for example, that each scientist is an expert in a very complex field that is not easily accessible to scientists of the same discipline. Say that it involves, for example, very specific experiments or calculations based on complex computational calculations that only very few people can make sense of. Therefore, there is no easy way for the other scientists to make themselves familiar with each other's work. Accordingly, each scientist has (for the sake of simplicity) a tracking reliability of 0.99, or any credence that suffices to call the scientists' beliefs on the proposition about which they are experts as knowledge.⁷⁸ In each scientist's case, he or she has a tracking reliability slightly below 0.5 on all the other 99 propositions. In the framework of judgement aggregation theory and List and Pettit's conditionalised probabilistic epistemological account, for such a case we can state the following. First, aggregating the judgements and letting each scientist judge on the proposition regarding her area of expertise, and not on the other propositions, will outperform simple majority rule. We are therefore confronted with aggregating the judgements of 100 propositions by means of dictatorship votes on each proposition, provided by different expert members. Second, as seen above, a premise-based procedure in which the 100

⁷⁸ I am using 0.99 as an example here, because there exists no clear threshold in Bayesian epistemology as to when a belief qualifies as knowledge. It is not necessarily the case that only beliefs that have a credence of 1 have to be conceived as knowledge. Also note that, assuming 1 as such a threshold may come with certain technical problems. Hence, to be on the safe side, I assume that 0.99 is close (enough) to being a belief that qualifies as knowledge in this specific case.

propositions are treated as premises and the conjunction, c , as a conclusion is advisable in this instance. Otherwise, the competence for the conclusion of each member would be incredibly low, or in fact practically zero. Consequently, we cannot rely on CJT to provide us with the benefit of a truth-tracking boost. Yet if an epistemology of true and reliable group beliefs is to carry any weight, it had better allow for such an instance to count as knowledge on a group level not only with respect to the premises but also, I assume, with respect to the resulting conclusion. In particular so regarding the conclusion, since it is supported by the judgements made on these premises and arrived at by means of an explicit or inexplicit premise-based procedure.

List and Pettit allow, in principle, for such instances to diverge from majority vote and yet to be an acceptable choice for tracking truth. It is not the case that there is a dogma that necessarily binds us to majority voting.⁷⁹ They instead argue that there exist various constellations of member profiles in which majority voting is less advisable. When a group should refer to a majority vote and when it makes sense not to do so will depend on whether, like in the scenario discussed, non-majoritarian procedures outperform majority vote. In other words, it will depend on whether the effect familiar to CJT can be outweighed by means of aggregation processes other than majority voting. Whether or not to opt for simple majority voting, or adaptations of it, depends on the member profile, as well as on the connections between the propositions on the agenda. Hence, there can be “epistemic gains from decentralization” that outperform the epistemic gains resulting from CJT in the case of majority voting.

Of course, there exist a lot of intermediate cases between a simple majority vote on the one hand and the example of our 100 scientists on the other. For example, we could alter the scientist scenario such that there are only 20 premises, p_1 - p_{20} , on which the 100 scientists are to judge and where for each proposition there are five scientists with a very high tracking reliability for determining its truth; say each of the five scientists has a

⁷⁹ Note that we are discussing judgements on propositions with truth value and these are concerned with facts. There indeed politically and morally ought to be such a dogma in the case of aggregating preferences in the context of policy making or other political decisions. I do not wish to suggest that what I discuss here could or should be applicable to political decision-making. I do recognise, however, that differentiating the cases discussed here from political decision-making processes may not be possible in many scenarios. It is highly important to clarify what differentiates one from the other, if such a differentiation is possible at all. At this point, however, I simply assume that there is such a distinction, even if only for the sake of argument.

tracking reliability (negative and positive) of 0.9 for the premise on which they are an expert. As in the example above, the scientists would each have a tracking reliability below 0.5 for all of the other 19 premises. In such a case, and given a premise-based aggregation procedure such as that outlined above, there would be 20 subgroups with five members each. Still, this would be a clear case in which simple majority rule would be outperformed by “decentralization”, that is to say, by letting the experts decide on the premises for which they have a particular expertise. In this scenario, however, there would still be certain (reduced) gains in relation to CJT because there are subgroups of five for each premise, each with an equally high tracking reliability.

While this altered scenario is easily determinable, there will be examples where it is not initially clear whether decentralization indeed outperforms the gains from CJT, which includes all of the members’ judgements on all of the propositions on the agenda. For example, if the experts’ tracking reliability were to be higher, but at differing degrees, it might still be the case that the scientist with the highest tracking reliability outperforms the gains that result from letting all five judge. With respect to decomposition, List and Pettit hold that the

key to choosing the right level of decomposition within a judgmental task lies in identifying, and designating as premises, those propositions on which individuals meet Condorcet’s competence condition. Any decomposition that is too fine-grained or too coarse-grained by this criterion may undermine the group’s ability to track the truth. (LP 2011, p. 100)

In light of the above, what should be noted for now is that there exists an essential dependency between the member profile, the group’s agenda and the chosen aggregation function. In order to determine how justified group belief is possible, all three aspects have to be given the relevant weight.

In summarising this chapter we are able to note that, first, List and Pettit adopt a Nozickian framework, which they interpret in a probabilistic manner. According to this proposal the knowledge challenge can be met if one is able to show that groups are capable of tracking the truth on a certain proposition reliably, where reliability depends on the positive and negative tracking reliability of the group members. Positive and negative tracking reliabilities are defined as (I) the conditional probability that the agent

judges that p , given that p is true and (II) the conditional probability that the agent does not judge that p , given that p is false.

Second, introducing a possibility result based on CJT makes for a strong, yet highly idealised case that indicates that a group can reliably track truths. CJT claims that a group's tracking reliability on a certain proposition p tends towards one with an increase in group size, given that: (a) the members defer to a majority vote; (b) each member has a greater tracking reliability than 0.5 and the same tracking reliability than all other members; and (c) their judgements have been formed independently, that is to say, for each judgement it holds that $\Pr(J^1|J^2)=\Pr(J^1)$, where J^1 denotes a member's judgement and J^2 another member's judgement towards p .

However, third, we have also seen that some of these conditions are unrealistic. In particular, the tracking reliabilities among members usually vary. A closer discussion shows, however, that the general effect described in CJT, call it a truth-tracking boost, can be saved at least in weaker forms in cases where several of these conditions are violated. In fact, we have seen that there can even exist various gains, not only from democratization due to CJT, but also from decomposition and decentralization in cases where no or only few gains from CJT are to be expected. Also, the present discussion shows that premise-based aggregation procedures in particular, or sequence priority procedures in general, are especially suited to tracking truths on a group level in many instances, at least in formal terms.

Finally, there exist three aspects that are co-dependent on one another in the course of ensuring truth-tracking in groups. These are member profiles, the group's agenda and the aggregation function that is or should be implemented.

Chapter 5

Externalism, Group Agents, and Two Forms of Reliabilism

After a fairly descriptive introduction to List and Pettit's account in general and their approach to truth-tracking in groups in particular, it is now necessary to clarify the lessons that can be drawn from the previous discussion and applied to the context of this thesis. In order to do so and to guide the discussion that follows, a critical summary and evaluation of the discussion so far is in order.

The general idea of List and Pettit is to conceive of group attitudes as being both dependent on and autonomous of member attitudes. The autonomy of a group attitude is warranted by the fact that group attitudes supervene holistically on member inputs. Furthermore, there are various gains, particularly epistemic ones, that result from a divergence from aggregation procedures and which would allow for direct reducibility to member attitudes. One challenge that a group agent and its enacting members face is that of ensuring rationality on a greater scale. Here it must be taken care that a group does not hold inconsistent, unrealisable or otherwise irrational attitudes or sets of attitudes and that the attitudes that a group holds are realisable and so on.

In addition to this general requirement, a connected but yet quite different challenge is that of ensuring truth tracking at the group level. Certain possibility results provide hints, indicating that groups may indeed be, under certain circumstances, capable of reliably tracking truths. To accomplish this it is essential for a group to dovetail three specific aspects relevant to the transition from a member profile to a group judgement that reliably tracks truth. What must be geared towards one another are (i) the particularities of the member profile regarding a proposition or set of propositions, (ii) the group's agenda of which the proposition or set of propositions in question are part, and (iii) the aggregation function. While List and Pettit discuss cases in which (i) can be altered given (iii),⁸⁰ in standard cases (iii) will be determined by (i) and (ii), that is to say, the choice of

⁸⁰ Straw vote procedures might be an example of (i) being, at least partly, determined by (iii). Also, in their discussion of discursive dilemmas, we have seen that sometimes (ii) is determined by (iii); this is so in that choosing a conclusion-based procedure, for example, leads to referring to some propositions on the group's agenda as premises and others as conclusions. Hence, there is a sense, albeit probably only a very

an aggregation function will depend on the member profile as well as on the logical connections between the proposition on a group's agenda. For the sake of simplicity, I discuss this standard case and hence assume that (iii) depends on (i) and (ii) in what follows.

There exists another lesson that we can learn from List and Pettit's discussion of a group's epistemic performance. List and Pettit make use of positive and negative truth tracking only in order to determine member judgement reliability. Given a certain profile of judgements and these judgements' positive and negative truth-tracking and indicating reliability, List and Pettit discuss the conditions under which these judgements can be "added up", that is to say, aggregated, such that the resultant belief is also reliable. As a result, the story they tell reads as follows. Imagine x members with determined positive and negative tracking and indicating reliabilities. Given this profile, this or that measure can be taken to ensure that, after aggregation, the resultant group belief is as close to 1 as possible. These suggestions concern the alignment of (i) to (iii) mentioned above. While, for instance, the requirement to pick a proper aggregation function – given the member profile – is present, List and Pettit's approach does not involve making this alignment an explicit epistemic requirement, but rather in treating it as a motivation in order to show, by means of possibility results, that groups can indeed be successful truth trackers.⁸¹ Hence they introduce "possibility results" and various examples that show which "choices" should be made given a particular situation, hinting at certain promising strategies as well as potential problems for arriving at a group belief that tends towards 1, or which at least outperforms the members' beliefs regarding its reliability.

While these discussions are insightful for discussing group knowledge, we may wonder whether possibility results and the adjunct discussions can be consulted for fleshing out an epistemological theory of justified (and true) group beliefs. If we turn back to Chapter 1, in which some basic features of externalism in general and process reliabilism in particular are introduced, we see that an externalistic framework primarily asks, *what has*

weak one, in which we can *make* some propositions premises for others. See List and Pettit (2011, esp. pp. 81ff).

⁸¹ This difference is difficile but significant. It amounts to the general approach for analysing group agency List and Pettit suggest. According to this approach, it is not so much the aim to point at the essential characteristics of a group, but to investigate ways in which a group can perform in a way such as to qualify as rational agent, successful truth tracker, responsible agent, and so on. (See List and Pettit 2011)

to be in place such that a given belief qualifies as knowledge?⁸² In the case of group belief the question would be, *what must have been in place such that a given group belief is reliable/justified?* However, List and Pettit approach it appears to be both compatible and divergent from such an approach. Their question is accompanied by an internalistic overtone. They ask, given the basis on which a belief is formed, *what should be done* so as to track truth? Their question regarding group belief is thus, *how should a group belief be aggregated such that, given a certain member profile and group agenda, the resultant belief is reliable/justified?* Hence, the debate includes the question of what steps an agent (group or members) has to take in order to warrant a group's successful epistemic performance. At the same time, however, List and Pettit adopt a purely externalistic stance with respect to justification, in which reliability (and hence justification) can be framed in a probabilistic version of Nozickian truth tracking. Accordingly, there exist rather internalistic requirements as well as externalistic ones and this might, at first glance, give us reason to worry about whether such a tension might be harmful. Hence this chapter aims to evaluate the degree to which and forms in which a group may depend on its members to contribute to the said group's epistemic success. Furthermore, it will be evaluated as to whether or not this might be compatible with a general externalistic framework, not only along the lines of List and Pettit's reliabilism, but also in light of a process reliabilist approach. Indeed, below I argue that asking this ultimate, rather internalistic question of what should be done to ensure a group's epistemic success is a possible move, even within externalistic epistemic frameworks; albeit only in the course of giving an answer to the former, genuinely externalistic, question and while caution is required in order to avoid toppling into internalism.⁸³ If anything, this insight should reassure us that externalism does constitute a promising framework for an epistemological analysis of group beliefs. Yet I will further show that while it is possible to manage this balancing act, there exists some reason to doubt whether, in combination with their rather internalistic interests, List and Pettit's (externalistically framed) possibility results are able to do so.

⁸² Or has a credence high enough to qualify as knowledge.

⁸³ I am grateful to Sebastian Kletzl for pointing out to me that there probably exists a confusion between externalism and internalism in List and Pettit's discussion. While I believe that the confusion is not as problematic as it might seem (see below), it is no doubt fundamental for this thesis to clarify this matter.

Moreover, this discussion will further lead to a closer examination of the specific kind of reliability that List and Pettit assume, in which the reliability of member judgements is not necessarily or directly bound to subjective degrees of belief, but merely to the degree to which member judgements are reliable given the proposition's truth or falsehood. As we shall see, while it is surely correct to proceed with an epistemological investigation of group belief on the assumption that judgements – as opposed to degrees of belief – are aggregated, this allows for the construction of cases in which a mere tracking-reliability account like List and Pettit's appears to fall short of satisfactory. Hence, we will see that the problem of List and Pettit's truth-tracking account does not so much lie in their utilisation of internalistic vocabulary, but in the particularities of their version of truth-tracking in groups.

5.1. Epistemic Externalism and Group Agents. Where the Problem Does Not Lie

One lesson that we can draw from List and Pettit's discussion of how groups can perform as epistemically successful truth-trackers is that in order to track truths a group needs to adjust its aggregation function depending on the member profile as well as the group's agenda. How should a group accomplish such an alignment? How should “identifying, and designating premises” (List and Pettit 2011, p. 100) and the identification of the distribution of probabilities regarding truth tracking among members be brought about? The most natural answer is that making sure that the proper aggregation function is chosen is something the members are responsible for. The reason for this is that a group itself, is, without its members' contributions, only a rudimentary agent, as seen in Chapter 3. Therefore, how else should any of this be accomplished, if not by means of members' actions?⁸⁴

We might suspect internalism to appear at this point, since we are no longer concerned with what has to be in place for a true belief to qualify as knowledge, but make

⁸⁴ Note that List and Pettit assume that members can (but need not) identify with their group, putting themselves in the group agent's shoes and thereby acting as well as thinking in its place. See List and Pettit (2001, esp. Ch. 9) However, and especially in such cases, it must be ensured that no requirements are placed on identifying members that cannot be met.

use of intentionalistic vocabulary, and require that the group agent somehow relates to its attitudes by means of its members doing so. We achieve this kind of relationship by suggesting what it is that the members have to be aware of, what they have to take care of and put into place in order to ensure a group's successful epistemic performance, and so on. Even the fact that members judge, which includes some kind of communication and/or action, might worry us. For instance, while in some contexts beliefs are treated more or less the same as judgements, in judgement aggregation the judgments of the members have to be provided to the group somehow. This can happen in various ways. In rather inexplicit settings, one possibility is to simply utter "p" in the course of a discussion. Another option is to click a box on a ballot, yet another to raise one's arm or say "Yes" or "No". Whatever it is, if we are talking about judgements being aggregated, there will be a (speech) act of some sort involved.⁸⁵

Yet I do not believe that we should be unsettled by all the possible types of demand placed on the members of a group that appear to be internalistic rather than externalistic requirements. Nonetheless, there are some serious worries regarding List and Pettit's model and the demands that probably follow from their discussion. Hence it is important to make clear which types of active involvement by members are permissible within an externalistic framework and which are not. In essence, there exist two ways in which members determine what happens on a group level. One way of achieving influence is by means of contributing a judgement to an aggregation process. The other kind of influence is less clearly defined. It amounts to various actions that ensure a group's epistemic success. Both of these dependencies of a group on its members will be particularly important for our future discussions. As we will see below, neither of these two ways in which members influence a group's epistemic performance need to be problematic for an externalistic analysis of justified group belief. I discuss this claim regarding both kinds of dependency below.

Member Judgements and Judgement Aggregation

How do things stand with respect to this first kind of dependency? List and Pettit are, it seems, quite right in claiming that we should not be distracted by the intentional

⁸⁵ See, for example, Searle (1969).

vocabulary used to describe the member level in the course of analysing a group's ontological and epistemic status. This is so because of the holistic supervenience relation between the individual level of the members and what happens there on the one hand, and the corporate level and what happens there on the other:

Under standard approaches to the individual subject, mind and agency supervene on what happens at the neuronal level. But the difficulty of predicting from a neuronal base what an agent does provides a justification for making sense of the agent in terms that abstract from the way its neuronal parts operate. Analogously, the difficulty of predicting from an individualistic base what a group agent does provides a justification for making sense of the group agent in terms that abstract from the way its members perform. It is true that in the case of a group agent we make sense of parts and whole in the same intentional vocabulary, whereas we do not use that vocabulary with reference to neurons. But this disanalogy does not undermine the fundamental similarity between the two cases. (List and Pettit 2011, p. 78)

If a group's attitudes supervene upon those of its members like a subject's intentional attitudes supervene upon her neuronal make-up and performances, then there is no problem in further claiming that a group agent's attitudes depend on the make-up and/or performance of its members. By adopting a perspective that focuses on the group agent rather than the members, we see that members' attitudes are to the group what neuronal processes are to an individual. This analogy seems to go together well with reliabilism, both List and Pettit's and Goldman's process reliabilisms. In the individual case, if these neuronal processes function reliably then our beliefs based on them will be accurate, if they malfunction or are unreliable, our beliefs will not reliably converge with the truth. In the case of a group, if the members perform or "function" well then the group will have reliable beliefs and accurate attitudes in general, if not, the group will encounter problems. The only discomfort within this analogy is that a group's "neuronal processes" include or are based on the *intentional attitudes* of the members and that there occurs an agential gap between individuals (the members) and the group. Rather than mere causal process like in case of our cognitive processes that are or are not reliable, in case of members providing reliability of a judgement on a group level, there will be actions, beliefs, and judgements involved. This takes some adjusting to, but it should not yet give rise to the rejection of an externalistic epistemological framework for analysing group beliefs. As we will see in a moment, issues regarding epistemic justification will precisely arise in discussing this very relationship. In particular, the question is how we are to

determine a member judgement's reliability. Is it enough to capture the rate with which an agent correctly judges that *p* (List and Pettit's truth-tracking account)? Or is there reason to resort to a process reliabilist framework in which reliability depends on the causal belief forming processes of the members, culminating in their judgements? In each case, however, the general claim holds that the fact that members judge in course of contributing to an aggregation process need not threaten an externalistic perspective on group beliefs. The question will not be whether or not these things are permissible in an externalistic framework, the question will rather be *how* and *to what extent* the justificatory status of member judgements contributes to a group belief's justification.

Ensuring a Group's Epistemic Success

The second kind of influence that goes from the member level to the group level includes intentional activities of various sorts to ensure a group agent's epistemic success. It seems unavoidable that members have to perform intentional actions and relate to their attitudes and those of the group in the course of choosing an aggregation function, and hence "choosing the right level of decomposition within a judgmental task" (ibid., p. 100) for example, hence aligning a group's member profile with its agenda and a proper aggregation function.⁸⁶ In general, it seems unavoidable that there exist members that ensure a group's rationality (see List and Pettit 2011, esp. pp. 104ff and 186ff).

Note that this case is different from the first kind of influence that member actions have on a group's epistemic status. In the first case, in which the relevant actions consist in judging a certain proposition in the course of an aggregation process, member judgements have a direct bearing on the reliability of the group's resultant belief. In the case of ensuring a group's epistemic success by aligning an aggregation procedure to a given member profile, for example, this *direct* relationship is not obvious and indeed not necessarily the case. They are, however, equally unproblematic. An example of this kind of dependence might help us to understand both the difference between these two kinds of member involvement and the degree to which they are unproblematic. Imagine you are really interested in figuring out whether *p*. As the member of a group that has *p* on its agenda, you get involved in preparing an aggregation process by, for example, making

⁸⁶ Moreover, it is even less avoidable in the case of group intentions and actions.

suggestions as to what the best-suited aggregation function might be. What is this activity – your activity – in relation to the group? It is in your personal interest to work out whether p and the group, and its capacity to aggregate member beliefs, are the right means to your ends. Of course, your actions will be relevant to the group’s epistemic success. However, in picking the proper aggregation function – based on personal motivation for the group’s success – the group is enabled to execute reliable processes that ensure its success in truth tracking. In such a case, what is the role played by your action in the group’s belief-forming process with respect to the resulting belief’s justification? Your action does not, I take it, form a part of those processes that determine a group’s reliability. It surely does not in List and Pettit’s version of truth tracking, since all that is required there is that members judge reliably. It is similarly unproblematic, however, for process reliabilism, in which the causal history of a belief’s emergence is subject to epistemic evaluation. To see why, an analogous example in the case of individuals will be of help. Imagine Liora, who is locked in a dark room. She has no access to the switch to turn on the light, since the switch is positioned outside the room. Yet she has to determine whether or not there is a chair standing in front of her. Since it is completely dark she is incapable of determining the chair’s presence or absence by means of visual perception. Now imagine that you are aware of her situation and, since you have access to the switch, you turn on the light.⁸⁷ This way, it becomes possible for her to perceive the chair that indeed stands in front of her. What is your action to the reliability of Liora’s belief-formation processes? Just as in the case of you picking an aggregation function, your action in the case of switching on the light to allow Liora to perceive the chair in front of her plays a constitutive rather than a justificatory role for the respective belief-formation processes. Surely your action is part of the (causal) history of Liora’s belief formation, as well as ensuring that a proper aggregation function is chosen is part of the causal history of a group’s belief formation. Yet, your role is in both cases, I take it, not part of the justifying history of these beliefs. In short, in both cases, your actions play a *constitutive* rather than *justificatory* role.

⁸⁷ Let us further assume that your turning on the light is not part of a shared or joint action of both Liora and yourself. However, there exist accounts of joint action, in which this example would probably still hold as an analogy (see, for example Bratman 1999).

Note that this is not the whole story. In addition to individually engaging in actions that promote a group's (epistemic and practical) goals, List and Pettit further allow that members may identify with the group agent, thereby making use of a "corporate" point of view.⁸⁸ In this case, members act, intend, believe, desire and so on, in accordance with the group's attitudes with the aim of ensuring a group's rationality or well-being for the sake of the group's rationality or well-being. They "assume a group identity", thereby adopting the group's viewpoint. This indeed poses an interesting case. What is a member's action for ensuring a group's epistemic success to the group in the case that the member is identified with the group agent?

In the above example, I concern myself with pursuing actions that ensure a group's rationality and/or epistemic success because of individual preferences and desires. In the case of being identified with a group agent, it is not my personal motivation that would be the reason as to why I act, but rather "we", the group, would seek to ensure "our", the group's, epistemic success. Accordingly, in such a scenario any action performed by a member would have to be attributed to the group. In this respect, the clear-cut distinction between constitutive and justificatory processes would surely fail to apply. Rather, such actions are the actions of the group. As thus they have to be viewed in the same manner in which individual actions are considered relevant or irrelevant in the course of an externalistic perspective of a belief's justification. While things might become more sensitive at this point, it would still be too hasty to dismiss groups as a

⁸⁸ Note that List and Pettit are, in fact, unsettlingly vague regarding the question of whether or not identifying with a group agent is a necessary requirement for group agency, and if, whether it is a necessary requirement for all members of a group to be so identified. For example, in Chapter 5 they indicate that identifying with a group agent is one of two possible routes in which a group's rational performance can be ensured. (See List and Pettit 2011, esp. pp. 124ff). An alternative way to ensure a group's rational performance would be the "organisational route" in which rationality is preserved by an organisational structure that warrants this. However, in the last chapter of their book, when dealing with identification with a group agent, they suggest that the organisational route is one incentive among others for identification with a group agent. Furthermore, while there exist passages in which it is clearly indicated that identifying with a group agent is a possible, and "often require[d]", yet not necessarily demanded for most or many group members (see esp. *ibid.*, p. 186), and, strictly speaking not necessarily required in each case, there are other passages that indicate that if identification with a group agent is not in place on the side of the members the group would "fail to perform appropriately as a group agent" (*ibid.*, p. 193). I fear that in discussing this issue we will have to live with this ambiguity. I will assume though, that identifying with a group agent is not a necessary condition for assisting a group in performing in an epistemically successful way.

proper object of externalistic epistemological accounts just yet.⁸⁹ Rather, it will depend on the very role we assign to these actions in the course of an account of justified group belief. As we shall see in the next chapter, merely stating that aggregation functions are conditionally reliable in relation to the member's profile is compatible with externalistic dogma, while capturing the importance of the relationship between member profile and aggregation function for matters of justification.

If identification with a group agent constitutes an interesting and different case for instances in which members act on the basis of their personal motivations, thereby ensuring a group's epistemic success, then the question arises as to whether being identified with a group agent may also constitute a special case regarding the first kind of dependency. Can we imagine cases in which a member who is identified with the group agent contributes to an aggregation of judgements, so that such a case somehow varies from the one discussed above? I believe that there is no clear answer to this question. One important note is that in the case of aggregating judgements regarding the desires or intentions of a group, being identified with the group agent and hence anticipating one's judgements as apart from what "I" want, and driven by what would be best for "us", the group, makes sense. However, things appear highly muddled in the case that such judgements regard propositions with a truth value. How are we to imagine an instance in which a group seeks to figure out the truths regarding a certain proposition and in which (at least some) members judge on the respective proposition in terms of "we" rather than "I"? Would such a judgement only include the reasons that a group might have for believing *p* to be true, just as in the case of desires, where a member identified with the group would give priority to the reasons the group would have (given its other desires and intentions) for having a particular desire? I do not see how to make sense of such scenarios. In particular, there exists a more general problem that concerns both the aggregation of desire-like states and beliefs. As seen in the previous chapters, aggregation procedures are key for List and Pettit's claim that groups are autonomous entities. In particular, aggregation functions are that which turn individual stuff into collective stuff.

⁸⁹ Indeed, there exist suggestions of how to include actions for increasing the likelihood of epistemic success in the course of a process reliabilist framework, at least for cases of individual actions. In particular, instances like the above might best be analysed along the lines of the justificatory role of epistemic virtuous actions. See Goldman (2002, esp. pp. 51ff) for one such proposal.

If an agent is identified with a group agent and while being identified as such judges that “p”, she will do so in corporate “we” terms or in another manner as an extended mind of the group.⁹⁰ But then her contribution to the aggregation function will not be that of an individual, but instead that of the group. Accordingly, the aggregation function will fail to aggregate individual stuff into collective stuff. We will encounter an objection that runs along similar lines in Chapter 7, where I discuss Goldman’s proposal concerning an intra-personal view of the relation between members and a group.⁹¹ At this point in the discussion it should merely be noted that contributing to an aggregation function process while being identified with the group agent provides a field for further investigations and may constitute a problem. But this is a general problem in List and Pettit’s account, and not so much a problem for an epistemological framework that is committed to externalism.

Moreover, it must be noted that since (probably) there is no necessity in being identified with the group agent in order to ensure a group’s epistemic success, this case is not of primary importance at this point, despite it being interesting. For now, it is important to note that – leaving aside identification with a group agent – there might be various types of action that are involved in a group’s epistemic success and performed by its members. These actions may even have the explicit goal of improving a group’s epistemic performance. Yet this does not necessarily threaten the idea that a group’s belief formation processes and the justification of those processes can be understood in externalistic terms.

In summary, there exist two types of a member’s involvement in the causal history of a group’s belief formation that need not be problematic from an externalistic perspective. Both such acts occur in the course of aggregating judgements for a certain proposition. This very relationship has already been identified as key to discussing group belief phenomena early on in this discussion. Actions of this sort are part of the history of a group belief’s reliability to the extent that they are part of conveying a certain belief. Or

⁹⁰ There is a discussion on „extended minds“ well-known from the philosophy of mind. I do not want to make a claim about members who are identified with a group agent serving literally as an extended mind off the group, albeit this claim does have some initial appeal. I am rather using this phrase here in an innocent sense. (See Clark and Chalmers 1998)

⁹¹ A comparable critique of List and Pettit can be found in (Rovane 2014).

actions are undertaken to increase the likelihood of a group's success in determining the truth of a proposition. Actions of this kind may play a constitutive role in the group's reliable belief formation, they do not, however, concern a group's justification with which it holds its belief (unless, perhaps, a member is identified with the group agent, thus acting as the group or its representative). Both kinds of member involvement are not, in principle, problematic to an externalistic, or more precisely reliabilist, perspective. Hence, this is not where the problem lies.

5.2. Truth Tracking. A Potential Problem

However, showing that tension is, in principle, absent from a member's intentional involvement in a group's epistemic concerns and externalism, does not yet mean that List and Pettit's reliabilist proposal is off the hook. A problem might still occur within the specific demands that are put on the members in List and Pettit's model, at least if we try to make the possibility results work for an epistemology of groups. It is one thing to allow for members to get involved in various matters concerning the group's epistemic success. It is another question, though, as to what demands are placed on members in order to succeed in such a task. There exist two potential problems with List and Pettit's probabilistic version of Nozick's truth-tracking account. First, if List and Pettit's tracking reliability account is to allow for members ensuring a group's tracking success, the members will need to be able to determine the reliability of their judgements, or those of their fellow group members. Second, it is unclear as to what exactly serves as a reference by which to determine a member's tracking reliabilities. To see why there might exist these problems, allow me to recap List and Pettit's account of truth tracking in judgement aggregations.

This requires that we briefly return to the conditions of positive and negative truth tracking and indicating reliability introduced above. According to this conception, what tracks the truth is not the specific degree of belief that a subject in fact possesses regarding p (with or without her knowledge of this), it rather denotes quite different probabilities from subjective degrees of belief. Truth-tracking reliability denotes the probability with which an agent will judge that p or $\neg p$ given that p is true or false.

Indicating reliability, on which the truth-tracking reliability of a certain group member depends, denotes the probability with which p is true given that the agent judges that p and false given that the agent judges that $\neg p$. The probabilities tracked by truth-tracking reliability regard judgements on a certain matter and, more precisely, they track how often an agent makes correct or incorrect judgements. As List and Pettit note,

each individual's judgment is akin to an independent coin toss, where one side of the coin, for example heads, corresponds to a correct judgment, which comes up with a probability of r , say 0.6, and the other side, tails, corresponds to an incorrect judgment, which comes up with a probability of $1-r$, say 0.4. (List and Pettit 2011, p. 89)

The two probabilities indicated in this quote are those that are formalised as $\Pr(J|T)$ and $\Pr(\neg J|T)$ in Chapter 4. The reason as to why List and Pettit obviously refrain from aggregating subjective degrees of belief is not so much that such degrees cannot be determined,⁹² but that there exists a specific gap between a subject's subjective degree of belief and the reasons as to why she judges p or $\neg p$. As we have seen already, judgements in the course of an aggregation process essentially involve some kind of act and hence a gap between subjective degrees of belief regarding p and an actual judgement on p . List and Pettit accommodate this gap in the course of their truth-tracking account. As a result, the merely rate with which an agent judges that p (or that $\neg p$) is measured, and this does not tell us anything about the epistemic reliability of the underlying belief per se. In general, rationality exceeds mere epistemic rationality and we may very well imagine cases in which maximizing one's utilities overwrites one's sincerity and performance as an *epistemically* rational agent.⁹³

Remember that in order to pick a proper aggregation function at least some members need to have access to the member profile. The member profile consists of a set of positive and negative truth-tracking and indicating reliabilities attached to member judgements. Therefore, in order to determine which aggregation function should be picked, at least some members have to have access to the member profile, and hence the distribution of reliability among the members with respect to their judgements. There are two potential problems that emerge at this point. First, in order to determine these

⁹² Yet this requirement would constitute a problem for externalistic models, should they be assumed. The point here is that List and Pettit do not choose this route. They do choose an equally problematic path though.

⁹³ See List and Pettit (2011, esp. pp. 104ff) for a discussion of this problem.

judgements' gains for truth tracking, the members would have to know about the proposition's truth or falsehood. To be clear, the rate with which an agent reliably judges that p might be determinable by means of past events in which p was the case and the agent did correctly judge that p , for example. For instance, let p (again) be that *there exists a surveillance camera in coffee shop C* and let us yet again consider Roberta, one of the robbers from Chapter 1 (see also ch. 5). We may determine her reliability for judging on the existence of a surveillance camera in coffee shop C on the basis of her past judgements regarding the camera's existence or the existence of cameras in coffee shops in general. While determining this will be difficult enough, the problem is that at least in one of these cases, the member determining the proper aggregation function will need to have had access to the truth of one of these judgements that Roberta made regarding cameras and coffee shops in the past.

Moreover, even if this is not a problem, there exists a second problem, which will sound familiar in light of some of the objections that Nozick's truth-tracking account has faced (see Nozick 1981). Recall Nozick's definition of knowledge and in particular his third and fourth condition according to which "(3) if p weren't true, S wouldn't believe that p ", and "(4) if p were true, S would believe that p " (ibid. pp. 172ff). One objection to Nozick's account is that it is not clear which possible worlds qualify as an indicator of reliability and which do not. Nozick himself raises this objection and as a consequence restricts the truth-tracking reliability of a certain belief to instances in which the same method has been used. Accordingly, his refined conditions read as follows (ibid, p. 179):

- (1') p is true
- (2') S believes, via method or way of coming to believe M, that p
- (3') If p weren't true and S were to use M to arrive at a belief whether (or not) p , then S wouldn't believe, via M, that p .
- (4') If p were true and S were to use M to arrive at a belief whether (or not) p , then S would believe, via M, that p .

Yet as Goldman (1983) points out, this restriction is unclear, perhaps necessarily so. It can be both too narrow and too broad, depending on how we conceive of "the same method". Is "the same method" a token experience, for instance, or the type of method like perception? Or does "the same method" depend on that for which it has been used?

In our case, does it concern the method of perceiving the very camera in question, or cameras in general, or rather only cameras hanging from ceilings? There exist, it seems, similar questions with respect to List and Pettit's probabilistic version of Nozick's truth-tracking account. What makes one instance of judging *p* an instance that serves as a reference point for determining a member's reliability of judging *p* while not another? In the case of Roberta, is her reliability in judging that there exists a surveillance camera in coffee shop *C* dependent on all her judgements regarding surveillance cameras? Or does it only include judgements regarding surveillance cameras hanging from ceilings? Does it merely involve judgements for surveillance cameras hanging from ceilings in coffee shops? Or is Roberta's reliability determinable only by reference to judgements regarding cameras hanging from the ceiling in coffee shop *C*?⁹⁴ Moreover, are only those judgements relevant that regard surveillance cameras hanging from ceilings in coffee shop *C* contributing to an aggregation of judgements in that group? More generally, what role does the fact play that the members judge on whether *p* in a specific context, namely in the course of an aggregation function?

These are all relevant questions, but I raise them only to set them aside. They show, however, that although the proposed account regarding the reliability of member judgements appears attractive at first, it is not at all clear how List and Pettit's truth-tracking account should be conceptualized once we dig deeper.

To be clear, it is not the aim of this chapter to show that reliability as List and Pettit assume it must be rejected. However, the previous discussion should show that there at least exists reason to doubt that such an endeavour was particularly promising or should be preferred over alternative accounts. As such, it is merely supposed to make room for alternatives. As we shall see in the next chapter (Chapter 6), there might be an alternative reliabilistic account on offer, namely Goldman's social process reliabilism. Before turning to this account however, I want to introduce one further potential problem of List and

⁹⁴ This latter restriction would be the closest to a judgement on the actual proposition in question, because it entails all three aspects: ceilings, cameras hanging from them and a specific coffee shop. Yet such a restriction would clearly be too narrow. Possibly, such a restriction would have to resort to some possible-world talk. In every (possible) instance of Roberta judging on whether there exists a surveillance camera hanging from coffee shop *C*'s ceiling: How often is Roberta's judgement positive (or negative) given that there is (or is not) such a camera? However, in this case, their account would clearly be susceptible to the original objections towards Nozick's possible-world account.

Pettit's truth-tracking account that will give further reason for me to address the question of how we are to conceive of the reliability of member inputs in the course of an aggregation process in different, yet reliabilist, terms.

5.3 Reliability of Belief vs. Reliability of Judgements. A Problem

So far, we have seen that epistemic externalism need not be afraid to include member actions and judgements in its account. I have identified two ways in which a member's action might have an epistemically significant influence on the group's successful epistemic performance. One way is that members judge on a certain proposition in the course of an aggregation procedure. As argued above, this kind of action is indeed relevant to a group's justification in the resulting corporate belief. Yet while its relevance has been stated, the extent to which an agent's judgement should be considered as influencing a group agent's justification has not yet been discussed.

Further, we have seen that there exist many ways in which member actions might contribute to a group agent's epistemic performance. Leaving aside the special case of members being identified with a group agent, I have argued that actions like aligning an aggregation procedure with an existing member profile, for instance, only play a constitutive role rather than a justificatory one. What counts from an externalistic perspective is that an aggregation function is dependent on the member profile. Even in instances where members act out of their commitment to the group agent, it still would not necessarily pose a challenge to externalism.

Yet we have also seen that there exist two potential problems with a probabilistic truth-tracking account like List and Pettit's. Both regard the second kind of involvement in a group's (epistemic) matters, that is, taking care of aligning a member profile with aggregation procedures. First there exists a problem in the "truth" part of truth tracking. It is likely that at some point the subject determining the member profile will have to have access to the truths of the proposition that is to be aggregated. Yet this might turn out to be difficult and, in some instances, perhaps impossible. Moreover, and in relation to this point, the second problem is that it remains opaque as to which kinds of

judgement are to serve as reference points for determining the truth-tracking reliability of members. However, both of these remarks pose challenges rather than objections to the kind of reliabilism that List and Pettit assume. Before turning to a discussion of a process reliabilistic framework for analysing group beliefs, I wish to add another challenge to this list.

Consider the two ways in which member actions have an influence on the epistemic performance of a group agent (again, let us leave aside issues regarding members who have identified with a group agent). While much of this chapter's discussion regards instances in which members act in ways to ensure a group agent's successful epistemic performance, it is now time to return to the former kind of influence, namely the kind of influence that stems from a member's contribution to aggregation procedures, which is indeed a kind of action that takes centre stage in dealing with the epistemic justification of group beliefs.

As mentioned earlier, operating with probabilities that merely capture the rate with which a member judges correctly or incorrectly on a certain proposition does not tell us anything about the actual belief a member has regarding *p*. As a consequence, a member's subjective degree of belief can come apart from her judgement. In a worst-case scenario, a member may strive to maximize her own utilities rather than judge sincerely, and hence may vote strategically and hence insincerely.⁹⁵ Could this pose a challenge to truth-tracking accounts like List and Pettit's? I believe so. The following example should help to pinpoint the problem, which will accompany us in later discussions regarding Goldman's process reliabilism as well.

Imagine a three-member advisory board for a school's new science curriculum.⁹⁶ Two of the members, Mary and Clair, believe in a "literal" understanding of the Bible and,

⁹⁵ Cf. List and Pettit 2001, esp. pp. 104ff.

⁹⁶ This case bears some resemblance with Lackey's CREATIONIST TEACHER example (I am grateful to the WFAP and in particular Sebastian Kletzl for pointing this out to me). See Lackey (2008, esp. pp. 48ff). Considering the discussion below, it might indeed be a particularly interesting thought to conceive of the Mary and Clair case as being the collective version of Stella, the creationist teacher. However, if so, the intuition that I suggest is present in Mary and Clair's case is quite the opposite of Lackey's intuition regarding the CREATIONIST TEACHER. Indeed, it appears that the case of Stella need not trigger the very intuition that Lackey promotes. Note that in order to acknowledge a resemblance between these cases, we would have to acknowledge that in cases concerning an aggregation of judgements testimony is (necessarily) involved. Yet it is too early at this point in my discussion to indicate this. However, see chapters 6 and 7 below.

accordingly, strongly believe in creationism.⁹⁷ However, when it comes to voting on whether creationism is a true theory (and as thus should be taught in science class) they judge that creationism is false. They do so with a negative truth-tracking reliability $\Pr(\neg J|\neg T)$ of, say, 0.8 and a positive truth-tracking reliability $\Pr(J|T)$ of, say, 0.2. This means that in cases where the members of the advisory board are to judge on whether creationism is false, given it is false, Mary and Clair reliably judge that creationism is false, while in cases where they were to judge on whether creationism is true, given it is true, they would be ill-performing.⁹⁸ Let us further consider the fact that creationism is highly unlikely to be true, say, $\Pr(T)=0.1$.⁹⁹

According to List and Pettit's model, Mary and Clair are highly reliable truth indicators, to degree 0.92, even though their positive tracking reliability is very low.¹⁰⁰ However, the reason behind this is not that their *beliefs* correspond with the truth. Rather, they judge correctly because their hope is that the religious studies teacher will convince the students that what they learn in biology is utterly false. Their reasoning is that if science classes were to mention creationism when discussing evolutionary theory then they would not come across as being stubborn. If the science teachers were to remain silent on the issue, however, the religious studies teacher could blame them for failing to be "objective".

In short, Mary and Clair adopt a very insidious plan. In line with their plan they judge that creationism is false (and should not be taught in science class) and they would do so whenever the vote is repeated. Their reasoning is probably rational on a practical level,

⁹⁷ I am grateful to Patrick Johannes Klug for making the most valuable suggestions as to how to improve this example, despite our disagreement over this example's implications.

⁹⁸ Above, I attest a very low positive tracking reliability for Mary and Clair. I do so because if Mary and Clair's negative tracking reliability is particularly high because of their strategic interests in the group's belief, they probably would deny that creationism were true even in the (highly) unlikely case that it were true. Because of the very unlikeliness of $\Pr(T)$, their low positive tracking reliability does not greatly influence their negative tracking reliability.

⁹⁹ Note that 0.1 is certainly already a higher degree than the actual probability of this theory's truth.

¹⁰⁰ Note that a feasible objection to this case could be that Mary and Clair could not have that high of a negative tracking reliability, since outside of the advisory board's meetings they would judge that creationism is true hence they would be bad at judging creationism to be false given it is false. Indeed, the previous discussion indicates that it is not clear whether only judgements in the course of aggregation processes are to be considered as references for determining reliability. Although quite far-fetched, there is no problem in principle in expanding this case so as to meet this ambiguity. Mary and Clair might be on a broader insidious mission, claiming to believe that creationism is false in any public context such as to infiltrate non-creationist circles, for example, for whatever reason.

but it is not *epistemically* justified. The third member, Rachel, also judges that creationism is false. She does so in accordance with her beliefs. Her tracking reliability both positive and negative is rather high, say, 0.85. Let us assume that she is not an expert on the issue but that she believes in what she was taught in biology and history classes at high school. Her judgements are therefore sincere and in accordance with her best epistemic abilities.

In such an instance, and considering Mary and Clair's contribution to the group judgement, would we really want to say that after aggregating the judgements properly the group reliably, hence *justifiedly*, believes that creationism is wrong? In List and Pettit's account we would have to consent that the group reliably believes that creationism is false.¹⁰¹ But does this not then point to a potential shortcoming in such frameworks? I believe that there is something out of place regarding justification on a group level in instances such as these due to the epistemic status of Mary and Clair's judgements in relation to their beliefs. When framing reliability as merely the rate with which an agent judges correctly, we are not able to explain what appears to be odd in such scenarios.

This provides, I believe, reason enough for us to re-evaluate the epistemic instruments that List and Pettit make use of in their discussion of how a group may be able to meet the "knowledge challenge". In particular, it provides another motivation for turning to process reliabilist frameworks. The question that arises in light of this example is the one identified in the course of discussing the extent to which judgements on a certain proposition in the course of an aggregation function can be analysed in reliabilist frameworks. It amounts to asking how, exactly, are we to depict the relationship between the member input and the way in which it contributes to a group's justification of the output belief.

Hence, in summary, this chapter has provided us with some transitional remarks from a discussion of List and Pettit's group account to an epistemological discussion of reliable group beliefs.

¹⁰¹ To be fair, List and Pettit spend a whole chapter on how the very gap between belief and judgement in the course of an aggregation can be closed. (See List and Pettit 2011, ch. 5) However, identifying this gap and making suggestions regarding how to close it does not yet say that such instances would not occur. From an epistemological point of view, it remains a challenge to rule out such instances for justifying a group belief.

First, it has been argued that there exist two ways in which member actions influence a group's epistemic performance. The first kind of influence regards member judgements in the course of an aggregation process. This very relationship has been identified as key to discussing issues of group beliefs as early in our discussion as Chapter 2. However, instead of aggregating beliefs, it turns out that what gets aggregated rather are judgements involving some kind of act.

One lesson drawn from the previous discussion of List and Pettit's group account is that it is an epistemically relevant requirement that an aggregation procedure is aligned with the specific member profile that serves as its input. Hence making sure that this alignment takes place constitutes the second way in which members may influence a group's epistemic performance. Both of these ways of influencing a group agent's epistemic performance are, contrary to the first impression and in principle, compatible with an externalistic framework.

While compatible in general, potential problems with each of the ways in which a member may influence a group agent's epistemic performance have been established with respect to List and Pettit's specific account of truth tracking.

With respect to the latter form of influence, two worries have been spelled out. First, in order for members to align the member profile with a proper aggregation function, the member(s) would have to have access to the truth of the proposition in question. Otherwise, it is difficult to imagine how the reliability of the various judgements of the members could be determined. Second, there exists a potential problem in determining what may constitute the reference points for determining the reliability of a member's judgement. This objection, List and Pettit's account owes its commitment to a Nozickian framework, in which similar objections apply.

With respect to the first kind in which members may influence a group's epistemic performance, namely by means of contributing to an aggregation procedure, one problem has been identified. Because aggregation procedures aggregate judgements on a certain proposition, p , rather than subjective degrees of belief, List and Pettit turn to an account of reliability that echoes this fact. While they are surely right to make this distinction, not including the reasons as to why one judges that the p I course of an epistemological investigation of justification of group beliefs turns out to be problematic. It allows us to

construct cases in which something (at this point not yet determinable) is odd regarding the justification of a group's belief in cases where it is based on members judging that p due to non-epistemic or otherwise epistemically bad reasons.

This discussion should suffice for us to now turn to alternative accounts of reliability and see whether they can do better. As we shall see in the next chapter, Goldman's social process reliabilism, and in particular his adaption of Sanford Goldberg's account of automatic justification transmission, might give us the epistemic instruments needed to avoid having to attest that everything is epistemically fine with group beliefs arrived at in such an odd manner as in the case of the school's advisory board.

PART III

Chapter 6

Adopting a Process Reliabilist Framework

The lesson drawn from the last chapter of the previous part is that ahead of us there still exist some challenges regarding an account of justified group belief. The first challenge lies in accommodating the dependencies that exist between aggregation functions, member profiles and a group's agenda.¹⁰² The second challenge rests in doing justice to the fact that there exists a gap between believing a certain proposition and passing judgement on it in the course of an aggregation process. In particular, given this gap, a truth-tracking account like List and Pettit's is incapable of accommodating cases where something appears odd with respect to reliability precisely because of this gap. The case of Mary and Clair illustrates one way in which this gap may lead to troublesome consequences for the particular version of reliability that List and Pettit assume. It is time now to turn to process reliabilism and see if this reliabilist proposal can do any better.

Chapter 1 contains a rough outline of process reliabilism. The reason for this discussion being rather sketchy is that it allows for a more fine-grained picture to emerge sequentially and in accordance with the idiosyncratic challenges of dealing with group belief. This demand still stands, but it is time now to introduce one specific and prominent account of process reliabilism and its extension for explaining group belief phenomena. As we shall see, this account does not solve all problems once and for all, but at least it provides some grounds for meeting the challenges identified in the last chapter. In a recently published paper, Alvin Goldman (2014) discusses the possibility of making process reliabilism work for group beliefs, in particular in light of the social ontological framework that List and Pettit provide.¹⁰³

One of the advantages that comes with adopting a process reliabilist framework lies in its commitment to being historical, in the sense that not only are the current states or circumstances of a belief's formation considered, but also the ways by which this current

¹⁰² While asking how determining the logical relationships between propositions on a group's agenda is an interesting question as well, I will neglect this requirement in what follows.

¹⁰³ Note that Goldman claims that, while he does adopt List and Pettit's theory, he believes that the epistemological questions he is concerned with can be clearly separated from ontological questions regarding groups. I very much doubt this claim, as we shall see in further discussions (see esp. Chapter 7).

state has come about. This allows for us to change our perspective with respect to the role that member beliefs play for the justificatory status of the group's belief. As seen above, List and Pettit's truth-tracking account acknowledges that the judgements provided by the members in the course of an aggregation procedure are not necessarily based on the subjective degrees of belief of these members. However, this acknowledgement comes with the price that their truth-tracking account does not differentiate between aggregating judgements that are reliable because they represent the members' beliefs and aggregating judgements that are reliable but based on epistemically non-permissible reasons or processes. The case of Mary and Clair introduced in the last chapter constitutes a case in point.

Assuming a gap between member beliefs and the reasons as to why they judge as they do is probably closer to the facts. As Goldman rightly points out (2014 p. 13, FN 1), talking of "judgements" in the course of aggregation processes points at utterances and hence speech acts rather than psychological states of belief.¹⁰⁴ So while the advantage of List and Pettit's formal account of truth-tracking is that it allows for such a gap, the resulting indifference in their account of truth-tracking is epistemically problematic in that it allows us to refer to member judgements as well as the resulting group judgements as reliable, even if the beliefs regarding the proposition they are to judge on are not justified and/or not even involved in the reasons as to why the member judges as she does.

Goldman's social process reliabilism does not provide a complete solution to this problem, since it does not close or even face the gap between judgement and belief in aggregation processes, but simply concerns itself with belief aggregation instead. However, as I show below, even though Goldman ignores the fact that members might pass judgement on reasons that are other than epistemic, an adaptation of his account might nonetheless provide the foundations for dealing with this problem. This would entail accommodating strategic voting as a possibility in groups by giving an account of this and explaining why a judgement (understood as a speech act, for example) based on strategic rather than epistemic reasons does not provide justification on a group level. I

¹⁰⁴ For Goldman, this difference is only worth a footnote in which he refers to this problem rather as a terminological one. I believe that the problem is more substantive than this and that it must be dealt with in the course of an epistemological analysis of group beliefs.

address this matter in the final chapter of this thesis. For now, let us postpone this discussion and proceed with an analysis of Goldman's social process reliabilism.

6.1. Goldman's Social Process Reliabilism

In order to identify and adopt the pieces of Goldman's theory that indeed constitute a step forward in group epistemology, it is necessary to take a step back from the discussion so far and ignore the difference between subjective degrees of beliefs and judgements in the course of an aggregation process and follow Goldman's social process reliabilism for a while. Accordingly, for now we will assume, without any reasons having been given, that member beliefs are aggregated at the group level in social process reliabilism. One of the key features of social process reliabilism is that, in addition to belief aggregation processes or functions, Goldman states that aggregating the justification of beliefs is a distinct matter. Keeping an aggregation of beliefs apart or distinct from what goes on at the level of justification carries the advantage of obtaining a better grip on how justification is able to transfer from the member to the group level in ways different from belief. In essence, beliefs come with a certain justification, that is to say, a certain reliability; this is based on the reliability of the processes that constitute the history of their formation. But the reliability of a certain belief is not or at least need not be accessible to the subject who holds the said belief. This is the externalist dogma. Accordingly, the story we are able to tell about the aggregation of belief states among members will vary from the story that we are able to tell with respect to the justification of both member beliefs and the resulting group beliefs.

Once again, an example proves helpful in demonstrating the core idea.¹⁰⁵ Imagine an adaptation of the Mary and Clair case from Chapter 5 for instances of belief formation. Although the Mary and Clair case illustrates a problem that stems from the gap between subjective degrees of beliefs and judgements in the course of an aggregation process, there exist cases related to Mary and Clair, at least with respect to the justification problem we face in such cases, on a group level. There exists a relation in that it too concerns cases

¹⁰⁵ Note that Goldman illustrates this point by means of his "British Museum Example", which seeks to explain more than this point and is therefore unnecessarily complex for illustrating the extent to which justification aggregation can come apart from belief aggregation (cf. Goldman 2014, pp. 15ff).

where an agent judges (or in this case believes) something for false reasons. In order to adapt the Mary and Clair case, let us focus on Rachel. Rachel judges in accordance with what she believes and this belief is aggregated properly via a suitable aggregation function. She believes creationism to be false and contributes this belief to an aggregation procedure. However, let us assume that Rachel believes so not because she believes in evolution, but because she is a member of a (different) religious cult that believes that humans were brought to earth three thousand years ago. Her beliefs neither include an Arc nor a God, but spaceships and aliens. For the sake of our argument, let us further assume that Mary and Clair each believe that creationism is false in a justified manner, for example, they are both evolutionary biologists and thus each have a highly reliable belief that creationism is bogus. All three believe that creationism is false and hence, given a majority vote, the group believes that creationism is false.

However, although the group believes that creationism is false unanimously, the justificatory status of the group's belief is not quite so monophonic. This is an instance in which the aggregation of belief comes apart from the justification that has been aggregated. Only the justification with which Mary and Clair believe creationism to be false will contribute to the group's justification in believing that creationism is false positively. Since Rachel's belief is not justified, because it is based on unjustified beliefs, this lack of justification also shapes the justificatory status of the group's belief, but negatively.

Splitting justification aggregations and belief aggregations appears to be a proper way to deal with parts of the problem that the original Mary and Clair case identified. In Chapter 5, I argued that focussing on the mere rate with which a member judges correctly that p deprives us from appreciating that a member's justification must somehow be related to the reasons she has for judging p . Goldman's proposal to keep belief aggregation apart from justification aggregation explains nicely that, and indeed how, the justification of a member's belief maps onto the justification of the group belief. It clearly rules out that in the adapted Mary and Clair case the group would be justified in believing that creationism is entirely false. For now, this verdict should suffice to attest that Goldman's social process reliabilism can meet at least these adapted versions of Mary-and-Clair-like cases. Before reintroducing the problem of subjective degrees of

belief and judgements in the course of an aggregation procedure, another problem that derives from Goldman's proposition must be addressed.

The question that arises at this point is: How can this approach be plausible? How, that is, can justification simply move from the member level to the group level? This question leads to the heart of Goldman's proposal and hints at two matters that are in need of further clarification. First, the very concept of epistemic justification needs to be clarified since it includes certain "tweaks" in order to accommodate the group case as opposed to the justification of the beliefs that individuals hold (see Goldman 2014). Second, how to overcome the agential gap between member beliefs and the aggregated resulting group belief must be addressed. Dealing with these two matters will further allow us to address the remaining challenges mentioned above, albeit in the course of critically reassessing Goldman's proposal.

Regarding the very concept of justification in groups, we are confronted with some idiosyncrasies regarding the justificatory status of group beliefs as opposed to comparable cases of individual beliefs. In individual cases there are two ways in which the processes that underlie the emergence of a belief can be reliable (see Goldman 2014, 1979; 2012a). A belief can be based on an unconditionally reliable process or on a conditionally reliable one. Unconditional reliability means that the reliability of a certain belief does not itself depend on the reliability of other doxastic states. Perception, for example, is unconditionally reliable in that (under normal circumstances) it produces significantly more true than false beliefs, independent of other doxastic states or processes. Reasoning, by contrast, constitutes an instance of conditional reliability. For example, if I believe q because I believe that p and that $p \rightarrow q$, then the reliability of my belief in q depends on the reliability with which I believe p , as well as $p \rightarrow q$. In such a case, the reliability of believing q depends on the justification of other beliefs, which serve as premises for that very belief.

Group beliefs are based on conditionally reliable processes (see Goldman 2014). These conditionally reliable processes are the aggregation functions. They produce reliable output beliefs only to the extent that the member profile is such that it allows for the chosen aggregation procedure to do so. This diverges from individual cases of belief reliability in at least two ways.

First, in the individual case there exists an inferential process from one or more beliefs and their reliability in arriving at a different belief, the reliability of which depends conditionally on the former. In the case of aggregating beliefs, however, one and the same belief about, say, p , is aggregated to a group belief regarding p . While this is uncommon, Goldman attests that this situation is nonetheless permissible in social process reliabilism (see *ibid.*).¹⁰⁶

Second, in the case of the reliability of individual beliefs, all of the processes that lead to that belief must have been justified. If p were not justified, then even if the belief that $p \rightarrow q$ were justified, believing q would not be justified. In the case of group beliefs, however, there might very well be unjustified member beliefs that are, together with justified member beliefs, aggregated to a group belief. Hence the justificatory status of a group belief must come in degrees and will have to allow for instances in which a group belief is justified, even though not all input beliefs are. Our altered Mary and Clair case, in which Rachel believes that creationism is false because she believes in a different, yet equally wrong ideology, may illustrate such a case. Her belief is not justified, but since Mary and Clair (now evolutionary biologists) justifiedly believe that creationism is false, Rachel's lack of justification does not render the resulting group belief unjustified, it merely decreases the degree to which it is justified. Note that, according to Goldman, the justification that is aggregated from the member level is an on/off matter, while on a group level the justification occurs in a matter of degree (see *ibid.*, esp. p. 28, Fn 11). Hence, while the belief aggregation process results in an unanimity judgement that creationism is false, there exists only a two-thirds majority regarding justification.

The resulting picture of group belief justification, including its dependency on the member profile, entails two assumptions (*ibid.*, p. 28 and 29):

(GJ) If a group belief in P is aggregated based on a profile of member attitudes toward P , then *ceteris paribus* the greater the proportion of members who justifiedly *believe* P , and the smaller the proportion of members who justifiedly *reject* P , the greater the group's level, or grade, of justifiedness in believing P .

¹⁰⁶ A lot of these "tweaks" as Goldman calls them are argued for in a rather lax manner, in the sense that the reason for allowing for them is merely that there is no good reason not to (see Goldman 2014).

And:

(R2G)(B) A group belief G that is generated by the operation of a belief-aggregation process Π is justified only if (and to the degree that) Π has high conditional reliability.

The first condition, (GJ), is concerned with the dependency of a group belief on its input beliefs, that is, member beliefs. The second condition, (R2G)(B), gives some indication as to the relationship between a belief-aggregation function and its justificatory dependence on the member profile. A belief aggregation function only yields reliable results dependent on the member profile. Remember that a belief aggregation function is not the same as an aggregation of justification. In fact, a justification aggregating function can vary from a belief aggregation function.

For example, a group might decide that a simple majority vote suffices to settle an issue, for example, whether or not p . Given that this majority is reached, a group will believe that p while the justificatory status will, however, depend on the justification of the input beliefs. It may thus be the case that a group believes that p , but does so in an unjustified manner, since not all of the beliefs that p , which form part of the majority of beliefs, are justified such that less than half may justifiably believe that p . In such a case, the group would believe that p , but unjustifiably so.¹⁰⁷

The combination of (GJ) and (R2G)(B) allows for the accommodation of both aspects of justification for group beliefs. First, it attests that the resultant group belief justification will grow with an increase in justified input beliefs. Furthermore, it is claimed that the reliability of an aggregation function depends conditionally on the member profile. Hence the requirement to do justice to the dependency of an aggregation function on the member profile is accommodated for. At least it is to the degree that a group's

¹⁰⁷ An appropriate example could be yet another adaption of the Mary and Clair case. Consider the altered version of Mary and Clair where the two are evolutionary biologists, but Rachel's belief is not justified. Now imagine that both Clair and Rachel were not justified in their belief that creationism is false because they believe in an equally problematic ideology as creationism that excludes creationism as a true theory. This would be a case that illustrates the aforementioned situation.

justification for a particular belief depends on the member input and on the conditional reliability of the chosen aggregation function given the member input.¹⁰⁸

While this much should suffice to motivate the core idea of epistemic justification on a group level, one open question remains to be addressed. This brief introduction to social process reliabilism allows us to better understand how the justification of member beliefs may determine a corresponding group belief in terms of process reliabilism. But we still lack an explanation of how such a transfer should be spelled out. How are we to show that justification can overcome the agential gap between members on the one hand and a group agent on the other?

6.2. Automatic Transfer of Justification. Two Models

There exists one familiar type of scenario in which we are confronted with such an agential gap in our epistemic lives. An agential gap occurs in cases where we rely on or learn from the words of others, that is to say, in cases of testimony. In essence there exist two camps within the debate about testimony, although there are a few accounts that probably take up an intermediate position.¹⁰⁹ The first camp comprises theories that claim reductionism (that is, inferentialism) in relation to testimony. The core idea here is that testimony is not one of our basic sources of knowledge, such as perception, inference, memory and, perhaps, intuition. Instead, its justification is mediated by one of the sources of knowledge and in particular by inference. In the reductionist understanding, a hearer must usually have some reason other than the mere fact that a speaker believes or utters that *p* to support her belief that *p* in light of the speaker's testimony.¹¹⁰

¹⁰⁸ For readers still concerned with the internalistic question of how members are to pick a certain aggregation function, simply stating that *there exists* a dependency on an aggregation function on the member profile might fail to satisfy. However, in an externalistic framework, this is all one will get with respect to analysing epistemic justification. Questions regarding how such aggregation functions can be determined in actual situations are a different matter. Hence this is a feasible way to circumvent the justificatory problem that comes with acting as a member such that a group's epistemic success is more likely, as discussed in Chapter 5.

¹⁰⁹ Such intermediate positions are Faulkner (2000); Moran (2005); Lackey (2008, 2010); McMyler (2011), for example.

¹¹⁰ Representatives of this line of thought include Hume (2008); E. Fricker (1995); Goldman (1999).

By contrast, non-reductionism (that is, non-inferentialism) denies this and claims that it *can* (but need not) be the case that we are justified in believing *p* merely on the basis of a speaker's utterance that *p*. Usually, the idea is that hearing a speaker utter *p* gives us some kind of default justification for believing that *p*, unless some defeater is present that would indicate otherwise. According to these accounts, testimony qualifies as a distinct source of justification and knowledge.¹¹¹

Can we learn something from these debates in order to make sense of how justification transfers from the member level to the group level? Goldman's answer to this is both yes and no. Yes, to the extent that he adopts a process reliabilist picture of automatic justification transfer transgressing agential gaps, thereby drawing on Sanford Goldberg's model of how justification transfers from one agent to another in cases of testimony. No, to the extent that Goldman attests that testimony understood as a scenario in which one subject testifies to another subject is ill-suited as a model by which to explain the (epistemic and ontological) relation between members and a group (see Goldman 2014, esp. pp. 19ff). Some more details regarding Goldberg's account and Goldman's adaption are thus due.

Goldberg construes his model of epistemic justification in testimony along the lines of what goes on in cases of memory and inference, that is to say, in cases of belief-dependent justification or what has been labelled "conditional reliability" earlier on in our discussion:

[I]t is standard for Process Reliabilism to regard both inference-based belief and memory-based belief as 'extended' processes, stretching back to include the process responsible for the production of the input-beliefs. It is true that, in the case of inference-based belief and memorial belief, the sort of temporal 'extension' that is involved does not extend to include processes that took place beyond the boundaries of the subject herself. But there is no reason in principle why the sort of extension that reliabilism endorses in cases of epistemic reliances *can't* do so. Testimony is a case in point. (Goldberg 2010, p. 67)

Since process reliabilism is historical in that it includes causal processes, there is no reason for us to deny the possibility that such processes might trespass the agential gap in cases like that of testimony. Memory might be especially illustrative to make a case in point. In

¹¹¹ Examples of non-reductionist theories include Reid (2011, esp. chap. VI, Sect. 24); Coady (1992); Burge (1993); Audi (2006); Kusch (2002); Sosa (2006); M. Fricker (2007); Goldberg (2010).

cases of believing something based on memory, process reliabilism would not only include processes of reassessing our formally formed belief, but also the very processes that led to the “original” belief. In the case of memory, accessing a memorial belief is something that happens within the subject, yet there is a gap between holding that belief at the time that it was formed, and reassessing it via memory at a later time. If overcoming such gaps within a subject by reassessment is possible and permissible, then why not across subjects? According to Goldberg, there is no difference in principle between the gap that occurs in cases of memory within a subject and the gap that occurs in cases of testimony between subjects that could allow a ruling out of the latter while allowing for the former. Moreover, once we allow for the idea that reliable processes can surmount agential gaps “Process Individualism” becomes questionable, which holds that “[f]or every subject S, all of the cognitive processes implicated in the formation or sustainment of S’s beliefs are cognitive processes that take place with S’s own mind/brain” (Goldberg 2010, p. 44).

Goldman adopts this model of allowing, in principle, epistemically relevant processes to cross an agential gap. However, he does not believe that such a trespassing is what happens (or can happen) in cases of testimony. This is probably a result of Goldman’s own commitment to reductionism about testimony. His argumentation holds that the “inter-personal model”, understood as a testimonial model along reductionist lines,¹¹² is ill-suited to explaining the very relationship between members and their group with respect to belief and justification transfer from the individual to the corporate level. This is why Goldman suggests that Goldberg’s idea of justification processes surmounting agential gaps should be adopted through an “intra-personal model”. According to this model the member belief profile, plus a conditionally reliable aggregation function, resembles reasoning from a group’s perspective, as opposed to resembling testimony.

Of course, adopting an intra-personal model does not say that what happens in groups directly resembles the reasoning of individuals. In the course of belief aggregation deductive reasoning from one proposition to another may occur, although it need not. Deduction on a group level may occur in premise-based aggregation procedures, for

¹¹² He especially refers to Jennifer Lackey’s reductionist justificationist account of testimony. See Lackey (2008 and 2010).

instance. In the course of such aggregations, the conclusion is drawn on a group level with respect to the premises that have been aggregated. But the aggregation is not, as such, of a kind that is familiar to reasoning processes such as induction or deduction. The step is not, for example, from p and $p \rightarrow q$ via deduction to concluding q , but from, for instance, p, p, p, p and $\neg p$ via aggregation function to p . However, this difference should not prevent us from viewing the transfer from member input beliefs to the group level as being akin to reasoning for social process reliabilism, as mentioned earlier.

In summary then, Goldman offers an account of how justified group belief can be made plausible in the course of a social process reliabilist framework. The key features of such a theory are, (1) to differentiate between belief aggregation and justification aggregation; (2) to make the justification of group beliefs dependent on both the member profile and the aggregation function, in which the latter is conditionally reliable on the former; (3) to allow the historic approach in which the relevant processes transgress the agential gap between members on the one hand and the group on the other; and (4) to conceive of the member input as being akin to the doxastic states on the basis of which individuals reason (the intra-personal view).

However, it is also clear that Goldman does not provide us with a satisfactory analysis regarding how we are to epistemically deal with the fact that members might not be truthful in their judgements. In fact, Goldman circumvents this issue by simply presupposing that the states that aggregation processes aggregate are belief states rather than judgements. As explained in the next chapter, this very assumption might backfire on his account, challenging, in particular, (4), while leaving (1) to (3) more or less intact.

The remainder of this thesis is concerned with identifying the key elements of the problem that Goldmanian social process reliabilism faces and offering some ideas for possible solutions.

Chapter 7

One Problem for Social Process Reliabilism and Two Solutions

Goldman's proposal of adopting Goldberg's account of justification transfer across agential borders for an intra-personal view is mainly based on the claim that the alternative inter-personal view is untenable. I first introduce the distinction that Goldman draws between these two accounts again and then argue that the intra-personal model is not a convincing answer to the question of how we should understand the relationship between members and a group.

As seen in the previous chapter, the intra-personal view suggests that the relationship between member belief inputs and the group belief output, determined by an aggregation of beliefs, should be seen as akin to reasoning processes. In this sense, arriving at a group belief based on member beliefs and an aggregation procedure is the group's way towards "reasoning".¹¹³ As Goldman himself notes, while this picture captures the "intimate relationship" (Goldman 2014, p. 20) between a group and its members, the main reason as to why Goldman adopts this model is that the alternative, the inter-personal view, is unconvincing in his eyes. A closer look at the inter-personal model and what speaks against it is thus in order.

While the intra-personal model takes reasoning as that to which it is akin, in essence the inter-personal model is construed as being analogous to instances of testimony. As mentioned above, there exist roughly two strands of testimonial account. Reductionism holds that testimony is not a basic source of knowledge because in addition to the testimony available to us, we further need some positive reasons in order to be justified in receiving the testimony and adopting the speaker's belief. Non-reductionism denies this requirement and argues that there are instances in which a speaker's utterance suffices for

¹¹³ It is briefly mentioned in Chapter 3 that groups cannot reason. This differentiates them from individual (human) agents and hence allows for the claim that there still exists a difference in quality between group agents and fully fledged (individual) agents. One essential ingredient of "reasoning" is that of taking a stance towards one's own attitude or set of attitudes. Yet in this context "reasoning" means something less demanding than what List and Pettit depict. It merely means inferring from one proposition to another, which we have seen that groups are capable of doing. Premise-based or more general sequential-priority procedures include this kind of "reasoning", as do procedures that are more closely connected to a richer concept of reasoning, as List and Pettit suggest; this includes adopting a stance towards one's own attitudes. Straw vote procedures are an example of the latter (see List and Pettit 2011, esp. pp. 59ff).

the adoption of the uttered belief, at least in absence of epistemic defeaters. Goldman construes the inter-personal view as a reductionist view of testimony. Moreover, he holds that a testimonial view on the relationship between a group and its members would be way too demanding and is thus ineligible.

The main reason as to why Goldman rejects the inter-personal view is this. In conveying one's belief as a speaker in a testimonial situation, the justificatory status of that belief is not automatically transferred along with it. Hence in addition to conveying one's belief, there is the further requirement of making transparent the reasons as to why one believes so. The more members a group has, the less feasible this demand becomes. Communication in groups would prove complicated and so a model that adopts the inter-personal view is unpromising (Goldman 2014, pp.19ff).

The second doubt that Goldman casts on the inter-personal model is that it is unclear who the receiver of a speaker's utterance could be. Clearly, it would have to be the group, but what exactly does this mean? Since the members are the very components of a group, and given that the testimonial acts would be votes or speech acts, how can a group receive these testimonies? Goldman asserts that in stating that the group is the receiver of the testimony, we mean that the members are. But if this is so, Goldman further attests, then each and every member would have to receive the testimony of each and every other member, in addition to their reasons as to why they believe what they believe (see *ibid*). If this were to be the case then forming a group belief would be so complex that it would, in essence, turn out to be impossible, at least in groups with more than just a few members.

All in all, Goldman concludes that "there are weighty reasons, therefore, to prefer the [intra-personal] model" (*ibid*, p. 20). But are these really such weighty reasons? A closer look at each of them should establish some clarity.

Let us begin with the first objection to the inter-personal model, which holds that a group's belief formation would be utterly complex because each member would have to convey the reasons for their belief in addition to that belief in order to establish justification for the resulting group belief. This claim is clearly only applicable for reductionist accounts of testimony. As seen above, Sanford Goldberg's non-reductionist model of testimony states that there *is* a transfer of justification across the agential gap between speaker and hearer, simply by way of the speaker's utterance. This is an

important point since it shows that one can advocate a process reliabilist theory while simultaneously rejecting reductionism about testimony. Hence nothing binds us to agreeing with Goldman that reductionism about testimony is true.

The second objection that Goldman makes is once again concerned with the complexity of a testimonial account for the member-group relationship. The problem is that it is unclear how a group might be the receiver of a testimony. According to Goldman, the group itself cannot be the receiver. This is plausible since a group itself does not have eyes, ears or brains unless these are “borrowed” from its members. Accordingly, Goldman concludes, each and every member would have to testify to each and every other member. Such an account would render group belief formation untenable, in particular in groups made up of a reasonable lot of members.

One feasible reply to this is to question one assumption that Goldman makes in the course of this objection, namely that *each and every member* will have to be the receiver of each and every member’s input. Considering the most complex and varying forms of how a group may be organised, it should be clear that (at least if we follow List and Pettit’s account) this requirement is by no means a necessary one. It could easily be argued that division of labour within a particular group might lead to only one or some members having to collect the testimonies of the members.¹¹⁴

This, however, is not the only possible reply. Recall that Chapter 5 left us with various challenges to group epistemology. In addition to the desideratum of meeting the specific form of dependency of a group’s chosen aggregation function on the member profile, another demand has been identified. According to this demand, it must be accounted for that there exists a difference between members’ subjective degrees of belief and the judgements that they make regarding whether or not that belief is true.¹¹⁵ Recall that List and Pettit’s model focuses on aggregation processes as the elements that turn individual stuff into corporate group stuff. Aggregation functions or procedures (as part of a group’s organisational structure) are the key elements to allow for a group’s autonomy with

¹¹⁴ Ideally, these members will have to be identified with the group agent (see List and Pettit 2011, esp. pp. 186ff).

¹¹⁵ The original Mary and Clair case was constructed on the basis of the gap between subjective degrees of belief and judgement. Ruling out, or at least somehow accounting for such cases, thus constitutes the third requirement outlined in Chapter 5.

respect to its members. This is so because in the course of an aggregation the resulting group attitude can come apart from the member attitudes in “surprising ways” (List and Pettit 2011, p. 69). Only by means of diverging from a direct accumulation of member judgements in form of simple majority rule can a group’s rationality be warranted.¹¹⁶ Accordingly, aggregation processes are key to both the “positive claim” applied to collectives and the negative claim that fosters group’s autonomy towards its members.

Now, given that the intra-personal model is true for a social process reliabilist picture based on List and Pettit’s account of group agency, how can this model circumvent the fact that judgements in the course of an aggregation procedure, that is, speech acts or other forms of communication, will be involved in aggregating attitudes from the member level to the corporate group level? In any case, intra-personal and inter-personal models alike, it will be the members’ takes on a certain proposition that will determine the group attitude. Given that these inputs are, as such, those of individuals and not those of the group, how might this work in the absence of an act of communication? Moreover, if an act of communication is involved, how might the intra-personal model circumvent the problem that Goldman identifies as being unique to the inter-personal model? To the extent communication is a problem for the inter-personal model, we see now that Goldman’s intra-personal model would be confronted with a similar problem, since it too cannot avoid communication among members and between members and a group.

How then do things stand with respect to Goldman’s third claim as to why the intra-personal view outperforms the inter-personal view? According to Goldman, the intra-personal model is apparently better suited to accounting for the intimate relationship that exists between a group and the members that comprise it. We might argue that Goldman surely has a point here. After all, there is indeed an intimate relationship between member and group. The group consists of its members and thus without any members, there would be no group.

However, the situation is not as simple as it may seem. A first reply might be constructed as follows. While Goldman is right to point to the intimate relationship between a group and its members, and while it is correct that this gives us reason to opt

¹¹⁶ While at the same time, as discussed in Chapter 3, rationality requirements must be violated.

for the intra-personal model, there exists at least one other reason that is equally intuitively appealing and which supports the inter-personal model instead.

It is important here to recall that List and Pettit take on a project that attempts to square the circle. Their aim is to argue that a group agent fundamentally depends on its members. Not only is it dependent on member attitudes as an input for forming group attitudes, but it also depends on members in order to achieve rational agency. At the same time however, it is claimed that a group is autonomous with respect to its members. This autonomy warrants that a group agent is indeed an entity in its own right, that is, a group mind, only less “scary”, since there is the dependence on side of the group with respect to the members. Clearly, the intra-personal model is supported by the fact that a group depends in the most fundamental of ways on its members. Yet it is equally clear that the inter-personal model is supported by the fact that a group is non-identical to its members and in fact personates an entity that is independent of them.¹¹⁷ There thus exists a standoff with respect to the intuitive support that these models enjoy, especially when assuming that the inter-personal model can be cashed out in non-reductionist terms regarding testimony.

To summarise, a closer investigation of Goldman’s discussion of both the intra-personal and inter-personal models reveals the absence of any good reason to reject the inter-personal model, at least if explained in non-reductionist terms regarding testimony. Goldman also fails to provide convincing positive arguments that speak in favour of adopting the intra-personal model. Rather, there exists a standoff between the intuitive support for the intra-personal model and the intuitive support for its rival.

Yet saying that there exists a standoff between the intra-personal and inter-personal models does not mean that there is no reason at all that could support one or the other. Indeed, when considering that Goldman’s social process reliabilism builds on List and Pettit’s account of group agency, there might be a reason to reject the intra-personal model.

In reaction to Goldman’s concern that an inter-personal model would be too demanding, since everyone would have to testify to everyone else in the group, I argued

¹¹⁷ Goldman briefly discusses the intuitive appeal of the inter-personal model with respect to its implications of non-identity between members and a group (see Goldman 2014, p. 19ff).

that speech acts or other forms of communication are equally required in an intra-personal model. Furthermore, I explained at the end of Chapter 5 and the beginning of Chapter 6, and noted repeatedly since then that if we are to take List and Pettit's account seriously then we will have to assume that there exists a gap between a member's subjective degree of belief regarding a proposition, *p*, and the member's actual judgement on that proposition in the course of an aggregation process.¹¹⁸ In fact, at the end of Chapter 5 I identify that accommodating this gap in some way presents a challenge to group epistemology if it is designed to build on List and Pettit's account of group agents. Moreover, as seen at the beginning of Chapter 6, Goldman appears simply to ignore this gap, stating that beliefs are aggregated and not judgements. In Chapter 6, I further suggested that a resolution of this matter be postponed in favour of focusing on the extent to which social process reliabilism is able to contribute important and informative insights in spite of the aforementioned concern. It is time now to readdress it.

Recall the two intuitions linked to the two models mentioned above regarding the close relationship between a group and its members on the one hand, and the fact that groups are (or can be) agents in their own right. If we assume a non-reductionist interpersonal model in which members testify to the group, and the member's justification for their beliefs automatically transfers to the group level, then we can accommodate both of the intuitions mentioned above. The members provide the necessary input for the group and they do so by means of judging (or uttering their take) on a particular proposition. In this scenario, individual stuff (the members' testimonies) is indeed the individual basis on which a group forms, via the aggregation of judgements, a corporate attitude. The intimate relationship between a group and its members can thus be accommodated by the non-reductionist part of testimony, although it of course requires further explanation.

However, if we adopt an intra-personal view, things appear to be less promising. The main problem is that in this model the very input belief would belong to the group level from the outset. In this scenario, something collective, namely the beliefs or judgements

¹¹⁸ They even devote a whole chapter to strategies regarding how to solve various challenges that result from this gap on a practical level. Note, however, that while they discuss this in greater detail the term "gap" is not used (see List and Pettit 2011, esp. pp. 104ff).

of the members *as part of the group*, are aggregated in order to constitute something collective, namely a group's attitude.¹¹⁹

In principle, there is nothing wrong with assuming that the judgements of those individuals constituting the group are part of the group's mental make-up,¹²⁰ however, such an assumption is problematic in relation to List and Pettit's theory. It threatens List and Pettit's whole account given the essential role that aggregation functions play as the very link between individual member and corporate group attitudes, which allows for claiming the autonomy of the group with respect to its members.

The inter-personal model does not appear to have this problem. Justification transfer in Goldberg's model works akin to the justification we get from premises in a reasoning process, but a testimonial belief is not per se a belief that is attributed to the hearer, it remains the belief of the speaker, which a hearer is justified to adopt. The difference between the individual instances of taking testimony and the corporate case is the following. While individuals accept a testimony unless defeaters are present and, in the case of multiple testimonies, reason their way to making up their minds given contradicting testimonies, a group's instrument for coping with testimonies is that of applying aggregation procedures. This would be yet another "tweak" of how to make process reliabilism fit for group phenomena. I do not see any reason for why this should not be feasible.

To summarise, the intra-personal model appears to have some difficulties. Not only is it difficult to picture how such a model might accommodate for the fact that judgements or speech acts, but not mere belief states, serve as the input for corporate output beliefs, but more importantly there is an ontological problem. For the intra-personal model to have any weight it needs to assume that member beliefs are part of the group's set of beliefs. If

¹¹⁹ Note that, in Chapter 5, I identify a comparable problem for aggregating judgements of members who are identified with the group agent.

¹²⁰ In what follows, I introduce an account to which the intrapersonal model can be applied that significantly diverges from List and Pettit's model of group agents. Also note that in Chapter 5 I argue that the member attitudes and actions are those upon which a group judgement supervenes, for example, individual mental states supervene upon an individual's neuronal processes. This is a different assumption from saying that member judgements are part of the set of mental states of a group, which, I believe, the intra-personal model would have to hold.

this is so, this constitutes a serious threat to perhaps the most essential claim that List and Pettit make in arguing for both the autonomy of groups and their dependency on their members. Aggregation processes, they claim, turn individual stuff into collective stuff. Furthermore, due to certain impossibility results such aggregations often do not allow direct reducibility to the member level, hence groups are viewed as autonomous with respect to their members. If member beliefs are part of the group's mind set however, the transformation from individual stuff to collective stuff dissolves because member beliefs are part of the group's mind set from the very start and as thus are essentially collective. Accordingly, the intra-personal model appears to be incompatible with List and Pettit's model of group agency, to which Goldman is actually committed.

Moreover, given the general commitments to non-summativism and process reliabilism (see Chapter 1) made in this thesis, it appears that there exist two possible paths that we might take from here. Both of these accounts attempt to save the general commitments of this thesis as laid out in Chapter 1. Hence, both suggestions assume that (a) non-summativism about groups is correct, which not only implies the assumption that groups are entities existing independent of their members but also that a groups attitudes are irreducible to the members' attitudes. Further, this thesis is still committed to (b) a process reliabilist epistemological framework. The two options introduces in the remainder of this thesis reflect these two commitments. One seeks to save the core idea of List and Pettit's aggregation account and attempts to do so by way of adopting a non-reductionist inter-personal view. Hence in this scenario, the social ontological underpinning remains fixed while the epistemological framework is changed such as to accommodate List and Pettit's group account. The second option seeks to save the intra-personal model, Goldman suggests, thereby laying to rest List and Pettit's group agents.

The remainder of this thesis is dedicated to introducing the potential candidates for each of these paths. Since both paths enter into uncharted territory, I restrict myself to only briefly outlining what I consider to be particularly promising cornerstones for such endeavours. This thesis is not, should not, and cannot be the text in which to develop new accounts. Even if it were, it is highly questionable as to whether its author would be able to provide a convincing alternative to what has been discussed so far. However, what this thesis can do is come to an end by pointing to potential lines of further inquiry.

7.1. Saving List and Pettit's Group Agents. A Non-Reductionist Inter-Personal View

While I believe that both paths are equally attractive, substituting Goldman's interpersonal model with a non-reductionist inter-personal model comes with significantly lower costs than abandoning List and Pettit's aggregation account altogether, which a revision of the intra-personal version would require. In fact, it might be a way of mending Goldman's suggested social process reliabilism based on List and Pettit's group agency account as opposed to being an alternative. The reason that it is treated as an alternative, however, is that Goldman would most certainly refuse to consent to a non-reductionist idea of testimony.¹²¹ Accordingly, from an epistemological perspective, this is, in fact, a big step.

How, then, might such a model unfold? Many of the cornerstones of a non-reductionist inter-personal model are outlined in this and the previous chapter. The core idea is that the relationship between members and a group in the course of a group's belief formation processes is one of testimonial belief formation. Accordingly, member inputs are treated as testimonies and the group's resultant belief as a testimonial belief. In non-reductionist accounts of testimony a hearer is entitled to accept a speaker's testimony, unless defeaters are present that indicate otherwise. However, this does not mean that we are epistemically entitled to believe whatever we are told right away. It simply holds that, in the absence of defeaters, we may adopt the uttered belief. Goldberg's testimony model, on which the assumed non-reductionist inter-personal view delineated above is based, holds that the justificatory status of the very belief uttered by the speaker transfers to the hearer in the case she adopts the speaker's belief. Applied to the group case, there is an automatic transfer of justification in cases where the members testify to the group. This avoids the problem of having to convey one's justification for believing *p* in addition to uttering that *p* that Goldman identifies as difficulty for the reductionist inter-personal model.

One key difference is that the standard case of testimony discussed in epistemology assumes that there is one speaker, one belief and one hearer who is or is not in the

¹²¹ Neither would Goldberg consent to using his account to explain how a group as a corporate entity can come to hold a belief (see esp. Goldberg 2010, pp. 186ff).

position to adopt the belief. In the case of group testimony, however, things look different, at least in those cases where neither dictatorship aggregation procedures are implemented nor unanimity is given. Probably the best understanding of what happens in a group's belief formation is that a group is confronted with various differing testimonies regarding a certain proposition. In the case of individuals being confronted with various (contradicting) testimonies, reasoning would most certainly kick in, since one testimony, regarding, say, p , would serve as the defeater of a testimony regarding $\neg p$ and vice versa. Groups, however, do not reason in this manner. Instead, judgements are collected and the resulting member profile is turned into a corporate judgement by means of an aggregation function (the reliability of which is conditional on the member profile). This is another adjustment that needs to be accepted if a process reliabilist approach is to be applied to group belief formation.¹²²

In order to get a better grip on this approach, it is useful to revisit the demands claimed to be essential for an epistemology of (List and Pettit's) groups at the end of Chapter 5. I have argued that there exist some demands for a group epistemology of List and Pettit's group agents. First, a workable group epistemology will have to capture the dependency that exists between a member profile and the chosen aggregation function (and a group's agenda). Second, it will have to be able to accommodate the fact that aggregation procedures aggregate judgements and not mere belief states. As thus, it has to accommodate the idea that there exists a gap between one's subjective degree of belief that p and one's judgement regarding p , since this gap allows for instances such as strategic voting. In context of this discussion, I have presented the case of Mary and Clair, the two creationists. A workable account of justified group belief, I have argued, must tell us what it is that is epistemically wrong with a group's resulting belief given Mary and Clair's judgements.

¹²² In addition, we have seen that there may be aggregation functions that more closely resemble reasoning in individuals. Straw vote procedures provide a case in point. Discussing in greater detail the various implications of especially functionally inexplicit aggregation functions within a non-reductionist inter-personal model would certainly be rewarding. For example, feedback allows not only for a group belief to hinge on member beliefs but also allows that a provisory group belief might be revised by member efforts to revise their own judgements by way of, for instance, joint deliberation. The result would be that transfer of justification would go into both directions. Unfortunately, such discussions cannot be provided herein.

Furthermore, we have seen that while social process reliabilism is capable of dealing with the first demand, and with an adaptation of the Mary and Clair case, Goldman simply ignores that members contribute judgements rather than beliefs and that the abovementioned gap exists, which allows for insincere voting behaviour. While in Chapter 6 this issue was placed on hold, it is now time to return to it and see to what extent a non-reductionist inter-personal view can accommodate it.

The previous discussion has already indicated that a non-reductionist inter-personal model is promising to meet these challenges. Remember List and Pettit's discussion of aggregation procedures. While their account is rather formal, we have seen that there are explicit, as well as functionally inexplicit, aggregation functions. We have also seen that there are various ways in which to implement or execute one and the same aggregation function. For example, a majority vote might be executed by using ballots, by raising one's arm or merely by saying "yes" or "no". There might even be a less obvious way to vote, for example, by defending some claim in discussion. All of these scenarios have the same function, namely contributing one's take on a certain attitude; in the epistemic case, to an aggregation procedure. Many, if not all of these ways to contribute can be viewed as instances of testifying.

Insincere judgements (as in the case of strategic voting), or conveying a belief one believes to be justified and true while it isn't so, are scenarios that all testimonial accounts have to be able to deal with. That there may exist a gap between a speaker's personal belief in *p* and his utterance regarding *p* is not, therefore, an inconceivable option. It is only so if we assume that the utterance regarding *p* is essentially identical to the very belief of the speaker with respect to *p*, which, unfortunately, Goldman's intra-personal view was forced to suggest. Goldberg's account of testimony might be particularly appealing at this point. To see how, let us reconsider the original Mary and Clair case and see how a social process reliabilist account along the lines of a Goldbergian inter-personal model tackles it.

In the original case (see ch. 5 above), Mary and Clair are members of a three-member advisory board concerned with a new curriculum for high-school biology classes. Both are passionate creationists and so believe that evolution is false. However, following a perfidious plan, they judge that creationism is false (and as thus should not be taught in

biology class). A third member, Rachel, also judges that creationism is false because she believes in the correctness of evolutionary theory. Therefore, the group unanimously believes that creationism is false. Yet I maintain that something is odd with respect to the group's justification of that belief.

A Goldbergian inter-personal view will most certainly allow us to claim that while the group arrives at a true belief, it bases its belief on unjustified (yet true) beliefs. Accordingly, the group is not justified in believing that creationism is false (although the belief is true). This is so because neither Mary's testimony nor Clair's testimony is justified, since neither are based on any corresponding reliable belief. Since the justificatory status of Mary and Clair's belief automatically passes to the group level, the claim that the group is also unjustified in believing that p becomes intelligible.

Therefore, to the extent that Goldberg's account can explain insincere or involuntarily false or unjustified testimonies, and to the extent that this lack of justification maps onto the level of the hearer, it is also able to explain the original Mary and Clair case, as well as comparable scenarios.

In light of the above, all of the demands identified in the previous chapters can be met. First, the conditional reliability of an aggregation function on the member profile is accounted for. This is a result of Goldman's discussion of how social process reliabilism can be applied to group belief. Second, that there exists a gap between a member's subjective degree of belief and her judgement in the course of an aggregation, which necessarily involves some kind of act, is accounted for as well. As a result, social process reliabilism along Goldbergian lines is fit to tackle cases such as the original Mary and Clair case.

Hence, this brief introduction to a non-reductionist inter-personal view on the epistemic relationship between members and a group should suffice to claim that indeed, such a perspective might be particularly appealing. However, there might exist one downside to this account. Remember that in Chapter 5, I identified two ways in which members can perform actions that are causally relevant for the reliability of a group belief. In course of this discussion I also briefly mentioned that according to List and Pettit, members might identify with the group agent, thereby adopting a group's viewpoint and thinking.

wishing, and acting as a group's representative. Essential to this kind of alignment with a group is that such a member will think and act by making use of first person plural terms. Hence, their reasoning will not involve "I, the member, believe that p", but rather "we, the group, believe that p". I also briefly mentioned that using first person plural terms in course of an aggregation procedure appear to problematic on both conceptual and ontological grounds. If what is aggregated are "we-attitudes", then how can an aggregation function still be the very process or mechanism by which individual stuff is turned into collective stuff?

At the beginning of this chapter, a similar problem with Goldman's proposal of an intra-personal view on the member-group relationship has been carved out. If members were identified with a group agent, it seems that the intra-personal model is much better suited to accommodate this kind of relationship. Indeed, the fact that Goldman's intra-personal view is confronted with quite similar problems should be evidence for its suitability. Yet, it is a highly problematic assumption that members might contribute to an aggregation function while being identified with a group agent. This tension is one that concerns List and Pettit's theory in general, it is not a genuine problem of a non-reductionist inter-personal model. Yet, such a model, if applied to List and Pettit's group accounts, will only be as successful as its social ontological underpinnings allow it to be. This thesis cannot provide an evaluation regarding List and Pettit's group account. However, its author is sceptical regarding this approaches success.

While this might constitute a satisfactory end to this thesis, a succinct outline of some cornerstones regarding how to save the intra-personal model is still in order. As opposed to the Godlbergian inter-personal approach introduced here, such a shift is more radical in that it must suggest a strong account of group agency that is more collective in spirit than List and Pettit's. Furthermore, it is questionable as to whether such an account of groups is compatible with the general idea of process reliabilism. Nevertheless, and as the final point of this thesis, allow me to introduce one such possible approach.

7.2. Saving the Intra-Personal Model. Rovanean Group Persons

Let us briefly call to memory the problems that the intra-personal view is confronted with if it is to be applied to List and Pettit's group agents. It is here where I argued that, essentially, two problems are faced by such an account. The first is that it is unable to tell us a plausible story of how belief states are aggregated rather than judgements, made public by some kind of act or by means of communication. That is to say, Goldman's intra-personal model does not explain how any *act* of conveying a certain belief can be part of it, which is particularly problematic considering that in aggregation processes belief aggregation will be mediated by members acting in the course of a speech act or otherwise to convey their judgement. The second problem is ontological. If the intra-personal model were accurate, we would have to picture the member beliefs as part of a group's set of mental states. In ontological terms, they would be collective in nature rather than individual independent contributions to judgement aggregation. This would mean that aggregation functions or procedures would no longer turn individual stuff into collective stuff. This would seriously challenge the core feature of List and Pettit's non-redundant realism, which depends on this function of aggregation functions for both their autonomy and dependency claims.

Accordingly, if an attempt to save the intra-personal model were to succeed, we would have to resort to a group account in which not one of these problems occurs. This means that, first, the chosen group account would have to provide a plausible story regarding how and why it is accurate, or permissible to state that beliefs are aggregated or contributed to the group level. Or otherwise it would have to show why no gap between beliefs and judgements conveyed via some act is possible between what a member believes and what she judges. Second, such an account would have to be a non-aggregative account of belief formation of groups. Or, alternatively, it would have to secure a group's autonomy on very different terms from those of List and Pettit, in particular in such a way so as to not depend on aggregation procedures in order to warrant a group's autonomy as well as dependency on its members. At the same time, such an account must be non-summative in the sense discussed in Chapter 1.

In the few remaining pages, one account that indeed meets all these requirements is introduced, namely Carol Rovane's group persons. While I do not believe that this is the

only account in the collective intentionality debate that might qualify as intra-personal, I believe that it provides the most developed account with respect to the epistemic dimension of group belief. In a recent paper, Rovane presents objections to List and Pettit's account which appear to be argued from a perspective that resembles the intra-personal discussed above (see Rovane 2014). This discussion, together with the account of group persons developed in her *Bounds of Agency* (1998), serves as a proper starting point. As we shall see, however, Rovane's account can be considered relatively radical in some respects. For instance, according to her picture of group personhood, aggregation functions are not rational instruments for arriving at a particular attitude of a group person. Moreover, there most certainly are various incompatibilities regarding the methodological grounds that she assumes and which social process reliabilism is committed to. After all, if Rovane's account is to be classified as either externalist or internalistic, it most certainly belongs to the latter. Nevertheless, I believe that some valuable insights can be drawn from an attempt to present some foundations regarding how to bring together an intra-personal view of social process reliabilism with group persons, although the resulting discussion might leave us with as many open questions as it provides answers.

Recall Figure 1 in Chapter 1, where the idea of non-summativism is outlined in a manner open to interpretation. According to this picture, there are members, or a group of people, and there is a group. Here a group is a separate entity, but it stands in a reciprocal relationship of dependency with its individual members; some of these relationships are constitutive and others are perhaps conceptual or practical. This representation is not perfectly suited to explaining the kind of non-summativism that Rovane develops. To see why an introduction to her account of group personhood will prove useful.

Her discussion of group persons stems from a wider concern, namely that of becoming clear about the concept of personal identity. That group persons can exist is compatible with her account of persons. A group person does not stand in any kind of relationship of dependence with its members, which are non-identical with it, but *it is* its members. At first glance, this might sound as if Rovane presents a distributive, reductionist account. In such an account a group would be nothing over and above its members and group

attitudes would derive directly from the members in the sense that, for instance, for a group to believe that *p* it would be sufficient that most of the members hold the belief that *p*. However, this would be a misinterpretation of Rovane's group persons. To see why, it is useful to briefly introduce her account of personhood in the case of (human) individuals and see how this account is applicable to the case of groups. This way, it should also be possible to pinpoint the extent to which Rovane's account allows for us to claim that a group person *is* its members and that at the same time non-summativism, and indeed non-reductionism, are both true for group persons.¹²³

Rovane rejects the idea that personal identity is conceptually (or ontologically) bound to what she refers to as "animal identity" (Rovane 1998). More precisely, she sides with John Locke on the following claims about personal identity (*ibid*, p. 14):

- (1) a person is something with a first person point of view;
- (2) the identity of a person consists in the unity and continuity of such a first person point of view; and
- (3) the first person point of view of an individual person need not coincide with an individual soul or an individual animal.

Claiming that personal identity neither necessarily coincides with the "individual animal" that we certainly are, nor with the "individual soul" (which most likely we do not possess), in this case amounts to roughly claiming that personal identity must be viewed as neither depending on some non-physical "soul" that hosts or constitutes our personal identity, nor simply that it is that which is bound to the physical unity of the human animal. Rovane diverges significantly from Locke's theory, however, in relation to what the "first-person point of view" is supposed to mean.¹²⁴ According to Locke, Rovane notes, the first-person point of view consists in "the phenomenological point of view of a unified consciousness." Rovane rejects this assumption in favour of defining a person's first-person point of view as a rational point of view (see, for example, *ibid* p. 19). What

¹²³ Note that List and Pettit assume a light version of group persons in that it is possible for members, on a voluntary basis and as a decision of the individual member, to be the group in the sense that I explain in what follows. However, they do not make this a necessary condition for membership in a group. Accordingly, the picture shown in Figure 1 applies to their account (see esp. List and Pettit 2011, pp. 186ff).

¹²⁴ However, at the same time Rovane believes that her reinterpretation of the "first-person point of view" as essentially being a rational point of view is Lockean in spirit (see Rovane 1998).

makes for a person, accordingly, is an “agent’s normative commitment to achieving overall rational unity within its rational point of view” (ibid., p. 130).

The concept of a rational point of view is essentially a normative one. It amounts to the very point from which a subject, or rather, a person, deliberates. In Rovane’s account, rationality covers rather commonsensical requirements like taking care of not holding inconsistent attitudes, having complete attitudes, that is, accepting and deliberating from the implications that a preference or belief might have, as well as ranking preferences, and so on. Moreover, overall rational unity amounts to not only exercising rational deliberation and action, but to achieving an overall rational point of view based on rational deliberation and action involving all-things-considered judgements.

Correspondingly, in Rovane’s normative analysis of personal identity, three criteria must be given for personhood. A person

- a. has a rational point of view
- b. is committed to achieving overall rational unity within its rational point of view
- c. is therefore susceptible to rational modes of influence, along with all of the forms of interpersonal engagement that rest on such influence. (Rovane 1998, p. 131)¹²⁵

Given that these criteria are sufficient to constitute personal identity, and considering that personal identity is separated from the phenomenological unity of consciousness as well as a human’s animal identity, two rather surprising consequences follow. First, one and the same single human being can have multiple personal identities, without any “core” person existing from which these identities are chosen. Second, one and the same person can be comprised of multiple human beings. Hence, group persons exist.

It should already be possible at this point to guess the extent to which Rovane’s claim is quite the opposite of summativist accounts. If personal identity can vary in one and the same “animal”, and one and the same personal identity can comprise many “animals”, then a group person will be one single unity, autonomous and irreducible with respect to any “members”, yet comprised of several animal identities that “host” that group person. Yet more needs to be said about certain features of these group persons in order to better

¹²⁵ Lowercase and no punctuation taken from the original.

understand how such an account avoids the three problems posed in relation to the intra-personal model.

One important remark regarding both individual (human) personal identities and group identities is that they consist in *achieving rational unity*. Emphasising the aspect of achievement allows us to see that there might be phenomena for which no full-blown *overall* rational unity is the goal, and hence no personhood, but were some degree of unity is still attempted. In particular, this allows us to see the ways in which group persons can come about.

According to Rovane, group persons can emerge as a special form of joint activity, where joint activity is defined as jointly striving for rational unity in relation to a joint goal. A brief look at joint activities and their connection with group persons, as well as the difference between group persons and joint activities, is thus of essence at this point.

It is the “extraordinary capacity of persons” (ibid, p. 137) that they are able to project themselves into another’s rational point of view. That is to say, it is possible for us to anticipate and elaborate upon the rational point of view, which does not belong to the very person doing the projecting. Yet in such cases there is a “basis person”, if you will, that projects herself into the rational space of someone else, while retaining the very perspective that one begins the projection from *as* the very perspective from which one projects. This is an important feature of persons in that it allows for their joint activities.¹²⁶ Two or more persons can engage in a joint activity, that is, in doing something together in order to reach a joint goal or a “perspective of common ends” (ibid p. 138). The capacity to project into the rational space of one’s companions in joint activities allows for various forms of coordination and anticipation in joint action and intention forming. However, this is *not* what goes on in the case of group persons. For cases of joint activity, it suffices that there are some goals to which those engaged in a joint activity agree upon. Each participant in a joint activity will attempt to reach rational unity with respect to the given joint goals. However, as opposed to instances of group persons, the aim is not (or need not initially be) to reach *overall rational unity*. Rather,

¹²⁶ Although perhaps controversial, I take it that when List and Pettit discuss identification with a group agent they have this type of capacity in mind. According to them, there exists an individual personal identity *from which* one chooses to identify with a group agent and to which one can return in the case that a group agent’s commitments conflict too much with one’s own (see List and Pettit 2011, pp. 186ff).

projecting oneself into the rational space of a joint activity does not change the fact that this projection is a means to the ends of the person that does the projecting. An example that Rovane herself chooses might help to exemplify the difference between joint activities and group persons (see Rovane 1998, p. 138f).

Imagine two students working on a philosophical problem together. One of the students is concerned with the premises and the other person is concerned with dealing with the conclusions and their implications. In this scenario, the two students may very well reason jointly. One relies essentially on the deliberation of the other and so their rational points of view are fundamentally interlinked. However, these two students engage in a joint activity rather than being a group person, or at least this is so if we add the following to the example. Imagine that one of the students is indifferent regarding the philosophical problem at stake. She only agreed to the joint activity because she needs to pass the philosophy class and in the course of this process has to solve the given philosophical problem. The other student, however, is completely committed to solving the problem, not for pragmatic reasons but because of a genuine interest in the problem's solution. Even if it were not required for class, she would still attempt to find the solution. This is not so for the other student. Accordingly, they each have different reasons as to why they engage in this joint activity and these different personal goals need not be aligned in the attempt to achieve rational unity. It is enough that they define a certain space of rationality, in which rational unity is required, like dividing labour, checking on the results, revising conclusions that appear to be odd, and so on. In essence, joint activities like these are characterised as being merely *means to a common end* for the individual (human) persons engaging in them. One wants to solve the problem as quickly and accurately as possible because she is interested in the result, while the other wants to make it easier for her by splitting the work. Consequently, there exists a rational space that they share and so a partial joint rational point of view; yet there is no overall rational unity between the two students. Still, as in the case of individual (human) persons, there is nothing that would not allow for instances in which *overall* rational unity is achieved by more than one human being in a joint manner. This is what characterises group persons. If an ensemble of human beings holds a first-person (plural) point of view characterised by being committed to achieving *overall* rational unity, the result will be a single personal

identity comprising many human constituents. Therefore, there is no point of view *from which* one is committed to rational unity in a different context than one's own (individual human) personal identity; rather, one is (a part of) the group person. A group person is thus defined as having a first-person (plural) point of view that is, as in the individual case, essentially concerned with achieving overall rational unity, and to which we can attribute a unified (group) consciousness. It is this rational point of view of the group that dictates the rational point of view of its members, not the other way around.¹²⁷ When there is a group person and when various (human-sized) persons merely engage in a joint activity is a matter of degree (see Rovane 1998).

At the beginning of this subchapter I claim that an intra-personal model would have to be able to handle two problems that the intra-personal model is faced with when applied to List and Pettit's group agents. The first is that an intra-personal model cannot simply state that belief states are aggregated rather than judgements or testimonies. The second problem is that if the intra-personal model were to be correct, the input member beliefs would be part of the group's mental states. As such, they would be collective rather than individual, which poses a serious threat to the role that aggregation functions play in turning individual attitudes into one corporate attitude.

Rovane's account provides a solution to both of these problems. With respect to the former, she claims that there is no reason for any human being who is (part of) a group person to contribute insincerely to the group's deliberations. It is, in fact, ruled out by the mere fact that a group person is committed to achieving overall rational unity:

[I]nsincerity would not be possible among the human constituents of a group person. By definition, such human beings would be bound by a commitment to achieving overall rational unity within the group. And such unity would make insincerity among them both unfeasible and unmotivated, in precisely the same way that it is unfeasible and unmotivated within the individual human-size persons we know. (Rovane 1998, p. 156)

The fact that insincerity is ruled out rests on the assumption that a person is defined by her first-person point of view, which essentially consists in being committed to achieving overall rational unity. Thus just as it is (conceptually) excluded that a single (human) person is insincere towards herself if she is to be rational, so it is conceptually impossible

¹²⁷ However, Rovane allows for this in cases where there might be rational points of view that remain in tact and which regard matters outside the group's rational unity. However, those rational points of view are "smaller than human size" (see Rovane 2014, p. 165).

that a group person is so without this being irrational.¹²⁸ However, this observation does not change the fact that in the case of group persons communication will take place. Rovane's discussion of sincerity and insincerity in communication acts is akin to this problem.

By reference to Gricean communication¹²⁹, Rovane parallels cases of communication within one person in the course of long-term intentions with what happens with respect to communication in joint activities and group persons. Communication is an important element for both long-term intentions in single (human-sized), persons as well as in group persons. The difference is merely that in one case the communication acts take place in a single animal identity, and in the other case between such identities. This difference, however, should be neglected. The result is that the difference between classically intra-personal and inter-personal relations blurs in that “*inter*personal relations are tighter, and hence more like *intra*personal relations, than we ordinarily suppose” (ibid, p. 161, emphasis in the original).

In light of the above then, there are two interconnected reasons as to why the first identified problem of the intra-personal model does not apply to Rovane's group persons. First, it is (conceptually) excluded that there is a difference between what one utters one believes and what one, in fact, believes. The reason for this is that the constituents (the “members”) of a group person are the group person in that they are committed to achieving the overall rational unity of their rational point of view. As a result, there exists no point of view from which they might be motivated to be insincere.¹³⁰ Second, the difference between utterances and belief states is blurred in that

¹²⁸ Admittedly, at this point Rovane's argumentation is more complex. She argues for this via an analysis of joint activities and long-term intentions, applying a model of Gricean communication to both and claiming that in group persons insincerity is impossible just as it is in joint activities or the individual (human-sized) case of persons (cf. Rovane 1998., esp. pp. 127ff). For our purposes however, the shortcut I outline herein will suffice.

¹²⁹ See Rovane (1998, esp. pp. 116ff).

¹³⁰ Note that List and Pettit also suggest ways in which to deal with insincerity and comparable problems that result from the gap between subjective belief and judgement. In particular, their discussion of how to meet the incentive compatibility desideratum, as well as their discussion of the possibility of identifying with a group agent is a point in case (see List and Pettit 2011, pp. 104ff). However, this discussion assumes that there exist members (likely the majority) who do not identify with a group in the sense required to be able to say that the same stated here with respect to their group agents. In fact, in their account it is not a necessary condition that all members are identified with the group agent. Hence the problem of having to account for the gap between judgement and subjective belief remains. Moreover, even

communication turns out to be something that also happens in individual (human-sized) persons in the course of long-term intentions for instance. Hence the inter-personal relationship of “member” communications to the group becomes less relevant.

What, then, about the second problem identified above? List and Pettit’s account of group agents essentially depends on the claim that groups can be autonomous agents because in the course of turning individual attitudes into corporate attitudes (via an aggregation procedure), member attitudes can come apart from the resulting group attitude. Their autonomy claim, as well as many of the ways in which a group depends on its members, is conceptually bound to aggregation functions and their role in group agency.

The first thing to note is that if Rovarian group persons were to aggregate judgements, it is probable that no problem would exist that is comparable to the problem that the intra-personal model faces when applied to List and Pettit’s group agents. First and foremost, this is so because every belief, utterance or communication process takes place in the “rational space” of the group person and is thus part of that group person’s mind set. Accordingly, there is no sensible distinction that is possible, or necessary, between individual member input and corporate group output. Apart from the general possibility of being able to aggregate judgements without risking losing the very link between the individual member level and the corporate group level, the aggregation of judgements is problematic in Rovarian group persons for differing reasons. In her discussion of List and Pettit’s group agents (see Rovane 2014), Rovane objects to List and Pettit’s account on the grounds introduced above. According to her, the aggregation of a majority of judgements does not amount to a rational procedure in a group person’s attempt to achieve a rational (and overall unified) point of view. While she does not exclude that voting is a potentially acceptable strategy (as some sort of last resort, perhaps), the reason as to why aggregation procedures are not particularly desirable in the case of seeking to arrive at a group attitude is that, if the perspective from which we are to reason as group persons is, in fact, the *group person’s* rational point of view, and given that

if it were to be required that all members are identified with a group agent, List and Pettit would run into the second problem identified above.

human individuals comprise the group person, achieving rationality is a matter of exercise deliberation rather than aggregation.

Note that List and Pettit would probably count deliberation as a form of judgement aggregation. However, Rovane insists on drawing a sharp distinction between these two procedures. The reason for this concerns the point of view from which one aggregates or deliberates. A group belief based on an aggregation is a function of input beliefs, whereas a belief resulting from deliberation is not a function of the beliefs that lead to it, but in fact a result of logical deliberative processes, just as in the case of (human) individuals who deliberate their way from various premises to a conclusion. Moreover, only in deliberations are the reasons as to why one arrives at certain conclusions available. In aggregation processes, these reasons remain unavailable.¹³¹ Only deliberation, and not aggregation, is, strictly speaking, rational. Hence deliberation – resembling a lot the example of the two students, only in course of a group person’s deliberation process – is to be preferred over aggregation in the course of a group person’s enterprise in achieving overall rational unity.

We thus see that one of the “tweaks” that Goldman suggests does not apply to Rovane’s group persons. Goldman suggests that, in the case of a group’s belief formation, something akin to inference processes is allowed and that here differing takes on the very proposition in question are added up so as to generate a corporate group judgement. The process of adding up member judgements is precisely what an aggregation procedure does. In the case of Rovane’s group persons, however, there is no need to allow for such processes to be part of the history of a group’s belief formation. The group is a unified person, despite being composed of many (human-sized) beings. While there still exists some kind of gap, this gap does not exist between one agent and another but instead an internal “fracturing” (cf. Rovane 1998) occurs through which the group person becomes many human-sized beings that reason as one unified body. As a result, the reasoning processes within a group can best be viewed as perfectly analogous to instances of individual reasoning. There is therefore no need to allow for aggregation procedures to be part of the history of a group’s belief formation.

¹³¹ The availability of reasons leading to a conclusion is, obviously, an internalistic requirement, just as achieving overall rational unity is. I discuss the extent to which an externalistic epistemological account may be applied to Rovane’s group persons below.

This brief discussion of Rovane's group persons should be enough to show why Rovanian group persons have to be considered as non-summative (and non-reductionist), while at the same time no apparent difference really exists between the individual level on the one side, and the corporate level on the other. It also shows why they can meet the challenges imposed on them in relation to the intra-personal view in social process reliabilism. A group person is, according to Rovane, one single personal identity composed of many human beings. Since one and the same human being can hold various personal identities, so they can hold (a manifestation of) a group personal identity. Hence, the autonomy claim holds since there is not even anything toward which a group person would have to be autonomous, because there are no member identities that would not coincide with the personal identity the group is. Belief formation in such groups (just as any attitude formation) starts out from a group's rational point of view and hence the components of a group person's (human being's) attempt to achieve overall rational unity *from this* very point of view. Second, accordingly there is no room for insincerity or strategic voting, in fact, the difference between inter-personal communication and intra-personal versions of it becomes negligible in that inter-personal communication resembles intra-personal communication more than we might have thought. Hence, there exists no gap between a member's subjective degree of belief and her judgement regarding that belief, even if communication is an essential feature of a group's deliberation process. This should provide a satisfactory reply to the first challenge that an intra-personal model has to face. Regarding the second challenge, we have seen that Rovanian group persons are not confronted with the problem of aggregation procedures failing to establish the link between the individual member level and the corporate group level. In fact, group persons are capable of deliberation just as individual (human) persons are. Aggregation procedures are thus a scarcely desirable method for arriving at a particular belief held by the group person, or must be seen as irrational (or perhaps non-rational or arational) altogether.

Up to this point, it might seem that the two approaches are perfectly compatible, at least regarding the cornerstones of Rovane's theory that have a direct bearing on the problems posed for the intra-personal model in the course of social process reliabilism. However, there is one worry that needs to be addressed before closing this discussion and

with it this thesis. This concern regards the general scope of Rovane's theory in connection to social process reliabilism. In particular, it regards the extent to which an obviously internalistic theory can be combined with a radically externalistic one. After all, I have introduced concepts and demands such as *achieving overall rational unity*; *committing to achieve such unity*; *explicit rules* that have to be followed for deliberation; and, in general, *first-person (singular and plural) point of view*, which *constitutes personhood* if (and only if) it is a *rational point of view* that demands to *being committed to achieving overall rational unity*.

It cannot be the purpose of this thesis to reconcile these two radically different ways of approaching epistemic issues. However, allow me to at least indicate one line of thought that might allow for such a reconciliation.

Obviously, the argument in Chapter 6 regarding the possibility of imposing an externalistic account of justified belief on List and Pettit's group agents does not apply in the case of Rovanian group persons. The "neuronal level" of a group person would simply amount to the "neuronal level" of the human brains that are part of it. Therefore, we may ask how, in losing an argument in support of an externalistic view on groups, an account of group persons, in which belief formation is viewed as being restricted by rational requirements and a *rational point of view*, to which a person is explicitly *committed* and which she *attempts to achieve*, can ever be compatible with justification hinged on reliable belief formation processes. Usually, epistemic externalism is viewed as existing in direct tension with internalism. Externalism even appears to be some internalists' worst nightmare.¹³²

But perhaps epistemic internalism, understood as a philosophical investigation into what it means for an agent to be justified in believing something, need not exist in a state of tension with epistemic externalism (and in particular reliabilism), which simply enquires into the reliable processes that exist and must be in place for a subject to hold not only a true belief but also knowledge. I believe that Foley (2004) has convincingly argued that epistemic internalism and externalism need not be viewed as rival accounts of

¹³² One does get this impression when reading BonJour. See BonJour (1980).

justification and knowledge.¹³³ Instead, they simply pose two very different questions and engage in very different projects. While Foley merely suggests that these two approaches might simply point at completely different philosophical projects rather than being in direct tension with one another, Anthony Booth (2010) has recently provided further support for Foley's compatibility claim. Given that this claim is tenable, Rovian group persons might indeed be compatible with an intra-personal version of social process reliabilism. After all, reliabilism neither presupposes certain restrictions regarding what it takes to be a person, nor regarding what it takes for an individual to be rational.¹³⁴ In the case of groups, reliabilism need not be interested in how group persons are constituted or even how they are sustained, it only matters whether they exercise certain processes that produce true rather than false beliefs in a reliable manner. Asking what it means for an agent to be rational and truth seeking is not of the reliabilist's concern.

Moreover, Goldberg's notion of automatic justification transfer in the case of testimony is adopted from his view concerning memory and inference, hence there is no trouble, it appears, in stating that in group persons a complex network of justification dependencies occurs within itself and thus among the human constituents that, via deliberation from the group person's point of view (that is, inference processes done by the group), result in a group person's all-things-considered belief. In fact, Goldberg has also "tightened" the relationship between intra-personal and inter-personal bonds with respect to justification, just as Rovane has done with respect to communication. As a result, such inference processes qualify as conditionally reliable, just as inference in general qualifies as conditionally reliable. The only difference is that they have to expand from one human being to another. In light of this, a group person can attempt to reach rationality and overall rational unity and knowledge without this having anything to do with questions concerning reliable belief formation and knowledge. Therefore, perhaps there is a chance for social process reliabilism building on Rovian group persons.

There is, however, one last point that needs to be addressed with respect to the compatibility of these two theories. Goldman's intra-personal model proves to be viable

¹³³ I am grateful, again, to Sebastian Kletzl for discussions regarding the internalism/externalism debate. This claim is most likely falls in line with his position on the issue.

¹³⁴ In fact, process reliabilism explicitly refrains from including rationality as a requirement for justification (see Goldman 2012a and see above, Chapter 1).

only if we adopt a radically collectivist perspective about groups (and perhaps one of the most radical accounts of collectivity in recent collective intentionality debate). In order to make this account work then, Goldman would have to depart from ontological and epistemic individualism in the most radical of ways. Not only is a human being fragmented, but personhood and hence rationality are not bound to human beings; they can instead take on different forms inside one human being. This is an outrageous claim for social epistemologists who are committed to metaphysical individualism, even if it allows for “group individuals”, as Goldman does. It is, however, an equally outrageous claim for more “moderate” anti-individualists such as Sanford Goldberg. Goldberg repeatedly emphasises that his aim is to question Process Individualism and that no stronger non-individualist claim is to follow from this. He explicitly distances himself from social epistemologies that support the idea that a groups exist as separate entities who can be the bearers of belief, knowledge or any attitude for that matter (cf. Goldman 1999, 2004, 2010). Thus from their perspective adopting the intra-personal model in order to make sense of whether a group can hold reliably formed beliefs that are true, and thus possess knowledge, comes at quite a price. So perhaps it would be the lesser of two evils to accept the inter-personal view as suggested in the previous subchapter. For those who welcome strong collectivist accounts, this chapter should reassure us that the only thing that stands between a reliabilist epistemology of group belief is the (potentially untenable?) commitment of these theories to certain aspects of ontological individualism.

Summary and Conclusion

This thesis is concerned with the possibility of justified group belief (and knowledge) being held by groups. Considering the many ways in which it is sensible to use the term “group”, and considering that there is more than one theory of knowledge (or justified belief) on offer, to begin the discussion some introductory remarks and preliminary assumptions were initially identified as necessary.

In light of the above, in the first part of this thesis two preliminary assumptions that should help to narrow down the field are introduced (Chapter 1). To accommodate for the fact that this thesis concerns itself as much with social ontological issues as it does with epistemology, one of these assumptions is regarded as the social ontological underpinning and the other the epistemological framework in which questions regarding the justification of group beliefs should be discussed. On the side of social ontology, this thesis commits itself to non-summativism about group entities, as well as non-reductionism regarding a group’s attitudes. On the side of epistemology, I have introduced process reliabilism as one of today’s standard accounts in epistemology. This framework has been selected not only because it represents a standard account in current epistemology, but also because it is, as of yet, the only “standard” epistemological framework where some groundwork regarding justified true belief on a group level has been provided in some detail. Following these two preliminary assumptions concerning ontology and epistemology, a working hypothesis regarding what might be the key question in such a discussion has been introduced (Chapter 2). It is argued that the very relationship between members of a group on the one hand, and the group on the other, understood in non-summativist terms as an entity in its own right, constitutes the best starting point for a discussion. In particular, both justification and how the justification of member beliefs regarding a proposition concerns the group belief that results from it are recognised as potentially the most relevant issues herein.

The second part of this thesis introduces one particularly well-discussed non-summativist account, namely List and Pettit’s group agents (Chapter 3). A detailed discussion of group rationality and List and Pettit’s holistic supervenience claim, as well as the very role that aggregation functions play in warranting both a group’s rationality and

autonomy, are provided. Moreover, List and Pettit's aggregative account of group attitudes includes suggestions concerning how groups can successfully track truths, and Chapter 4 provides a detailed overview of this discussion. List and Pettit assume that groups can track truths. The upshot of their fairly technical discussion is this. First, the member judgements on a particular proposition, p , need to be reliable, where reliability is measured in probabilistic terms, building on Nozick's truth-tracking account of knowledge. Second, there exist various aggregation functions that promise a group agent's epistemic success, *given* a particular member profile. In this sense, there is no "one-size-fits-all solution" since whether or not to apply majority rule, for example, will depend on the particularities of a given member profile. Subsequent to this fairly descriptive discussion, Chapter 6 constitutes a transition chapter in which List and Pettit's account is evaluated and certain problems that their version of truth-tracking potentially has to face are addressed. Two corresponding types of actions are associated with each of these two features of successful truth-tracking in groups.

First, it is noted that judgements in the course of an aggregation function must be conceived as (speech) acts and do not necessarily represent a member's actual subjective degree of belief. List and Pettit's reliability model does accommodate this fact in that merely the rate with which a member judges that p is considered relevant for a judgement's reliability. Nonetheless, this framework is not free of potential problems. Notably, focusing on the rate with which judgements are correctly passed in the course of an aggregation process allows for constructing counter examples in which something appears to be at odds with a resulting group belief given that member judgements are reliable according to List and Pettit, but still appear not to be.

Second, the demand to align member profile and aggregation functions also involves member actions. In order to achieve such an alignment, it is argued that at a certain point the respective members would have to have access to the truths of the proposition in question or to comparable propositions that lead to determining a group member's judgements as reliable or unreliable. However, there exists a more general worry regarding reliability accounts like the one List and Pettit assume. Their adaption of a Nozickian model of reliabilism (which they spell out in probabilistic terms) comes with a lack of clarity with respect to what kind serves as a reference judgment exactly in

determining the reliability of a member's judgement. A similar kind of problem exists for Nozickian truth-tracking accounts with respect to which possible worlds should be considered as constituting a reference in order to determine tracking reliability.

While action is involved in both cases of a member's influence on a group agent's epistemic performance, I further argue that this does not necessarily pose a threat to a general commitment to an externalistic epistemological framework and also explain why this is so. Hence, in light of the particular worries identified with respect to List and Pettit's truth-tracking account, Chapter 6 introduces Alvin Goldman's proposal of how process reliabilism can be adapted to accommodate group belief.

Yet this account has also proven to be somewhat unsatisfying. Goldman's proposal provides many interesting and perhaps even ground-breaking observations and suggestions. Moreover, it manages to accommodate the reciprocal relationship between member profile and aggregation function. Yet it fails to accommodate, or even appreciate, the gap that clearly exists between a group member's belief regarding a certain proposition and her actual judgement in the course of an aggregation process. Instead, Goldman simply claims that social process reliabilism (and in a sense List and Pettit's aggregation account) is concerned with mere doxastic beliefs rather than judgements in the sense of utterances or other forms of acts (like checking a box on a ballot).

As a consequence, Goldman's proposal regarding how a transfer of justification should be thought of becomes problematic, as does the relationship between members and a group regarding belief accumulation in more general terms. He suggests that the very relationship between a group and its members' attitudes should be thought of along the lines of what he calls the "intra-personal view". According to this view, member input beliefs are available to the group, just as premises are to a subject in inference processes. At this point, however, Goldman's assumption that beliefs, rather than judgements or utterances, are aggregated backfires and the intra-personal model not only appears to be inappropriate in its own right, but also demonstrates a problematic consequence for List and Pettit's account of group agency. If the intra-personal model were correct, member beliefs would be part of the group's mind set. This is highly problematic on the grounds that if member input beliefs were to be part of a group's mind set then an aggregation function would fail to be that which turns individual stuff into collective stuff, as List and

Pettit claim. While this would probably not constitute a problem for other group accounts, it is a problem for List and Pettit's since aggregation functions serve as the interface between the member and the group levels, thereby warranting both a group's rationality and autonomy.

Accordingly, I conclude that Goldman's intra-personal model is ill-suited to List and Pettit's account of group agency. Chapter 7, the final chapter of this thesis, identifies two possible solutions to this problem. One is to depart from the suggested epistemological framework while maintaining the social ontological underpinning (List and Pettit's group agents) in a fixed state. The second solution is to keep the epistemological framework in a fixed state (hence Goldman's social process reliabilism) while changing the social ontological underpinning by suggesting a different account of non-summativism about groups that allows for an intra-personal view.

In this first scenario, we have seen that adopting Goldberg's non-reductionist account of testimony (which Goldman refuses) and his account of justification transfer across agents (which Goldman adopts for his intra-personal model) allows us to give an account of how group belief can be reliably produced. Moreover, it is argued that this fits nicely with List and Pettit's group agency account and is indeed capable of accommodating the three problems identified in the course of a discussion of List and Pettit's account of truth-tracking in groups. Yet in such a scenario reductionism about testimony must be dropped and with it Process Individualism, which claims that reliable processes relevant to epistemic justification have to take place in a single subject. We have also seen that adopting Goldberg's model of non-reductionism about testimony cannot be constrained to apply to group phenomena only. Rather, it would have to be a general claim about testimony (as it is in Goldberg's theory).

In the second scenario, a suggestion as to what a group account would have to look like if the intra-personal view of social process reliabilism were to be adopted has been spelled out. These criteria alone already show that if the intra-personal view were to be adopted, a rather radical account on the side of the social ontological underpinnings would have to be assumed. One such account is Rovane's group persons. After briefly introducing the cornerstones of her account, the extent to which this account might prove compatible has been discussed in brief.

So what is the upshot of all of this? Why spend more than one hundred pages carving out which process reliabilist framework might be best suited for List and Pettit's group agents? The explanation for this is as follows.

Recall the list of features Goldman identifies as criteria for "standard" or "real" epistemology presented in the introduction. Those features were:

(A) The epistemic agents of traditional epistemology are exclusively individuals. (B) Epistemology focuses on the study of epistemic evaluation or normativity, represented by such evaluative concepts as justifiedness, rationality, and knowledge. Traditional epistemology asks such questions as how individuals can acquire knowledge and maintain justified or rational credal states. (C) Traditional epistemology assumes that the normative standards of rationality and justifiedness are not merely conventional or relativistic, but have some sort of objective validity. (D) The central notions of epistemic attainment – knowledge and justification, for example – either entail truth or are closely linked to truth. A known proposition must be true; and justified beliefs, according to many mainstream views, are in some sense likely to be true. (E) Truth is assumed to be an objective, largely mind-independent, affair. (F) Traditional epistemology takes its central business to be the critical examination of doxastic 'decision-making' (DDM): adopting, retaining, or revising one's beliefs and other doxastic attitudes. (Goldman 2010, p. 2)

This surely represents a proper picture of standard epistemology's self-conception. However, there are many aspects that might deserve greater scrutiny, at least in a revisionist spirit.¹³⁵ In particular, (C), (D) and (E) might appear unconvincing from a rather social perspective. Yet there is another assumption that can be challenged, namely (A), and indeed this assumption is targeted in this thesis.

List and Pettit's account of group agency is particularly attractive for philosophers who are willing to go along with some aspects of collectivist thought while seeking to keep intact various parts of the individualist dogma, as well as the idea that justification or truths are not a social matter. Groups may exist more or less "really" within such a framework and they may even gain a certain degree of autonomy. But in the end it still boils down to (human) individuals who need not (but may) chose to identify with a group. In this scenario, individual autonomy over one's own attitudes and over the group one may or may not choose to be a member of is warranted. It is warranted at least in a loose way by means of aggregations of judgements, even though the resulting judgements

¹³⁵ In fact, claiming that a revisionist approach is disqualified from epistemology is a daring and most certainly dogmatic claim.

may come apart from the set of member judgements. In the end, however, the members are still the masters of their mental set-up, their beliefs, desires, hopes and goals, and, in particular, the rationality and justification with which they hold their beliefs is not a social matter. We are able to retain all of our relatively comfortable theories regarding practical rational, belief justification and epistemic rationality. There is, in the end, no social influence that could “manipulate” our purely individual rationality as group members in any relevant sense.¹³⁶

Goldman believes, it seems, that he could stretch (A) so as to cover group phenomena, probably by insisting on the intimate relationship between members and a group or by conceiving of group entities as (non-human) individuals or by claiming both. But as I hope to have shown, such a theory about groups, of which List and Pettit’s is surely the most elaborate available at the moment, is not compatible with Goldman’s process reliabilist picture unless more concessions are made towards a more social picture of group beliefs. The discussion of social process reliabilism and List and Pettit’s theory serves as a token discussion. It is supposed to show, if indeed successful, that as soon as we allow for groups to play any epistemically interesting role, only minimal concessions to collectivism will not suffice. In the case demonstrated here, social process reliabilism will either have to adopt Goldberg’s account of non-reductionism about testimony. In this case, justification will overcome an agential gap in instances of mere testimony between individuals, as well as in the course of members testifying towards the group agent, which poses a challenge not only to (A) but probably also to certain aspects of (B). Or it will have to commit itself to even more radical claims about the social dimension of knowledge in which rationality, and probably justification, must be viewed as relative to the personal identity that engages in achieving rational unity. This, in addition to challenging (A) in a radical way, also poses a problem to, especially, (B) and, probably, (C).¹³⁷ From a perspective sympathetic towards collectivism, showing that at a minimum Process Individualism must be dropped may not appear to be a particularly unsettling claim, and indeed may even appear as a rather weak conclusion. However, it is unsettling

¹³⁶ Unless we choose to identify with a group agent, see List and Pettit (2011, esp. ch.9). However, there still exists the individual level of identity through which we place ourselves into the shoes of a group agent.

¹³⁷ Rovane indeed defends some version of relativism, see Rovane (2011).

for social process reliabilists like Goldman. Equally, that Goldberg's irreducibility account of testimony at least fits perfectly with such moderate accounts as List and Pettit's group agents, and that as thus the account *allows for*, in principle, the building of a group epistemology around it, is probably not good news for him too. All that prevents such an account from being applied to groups is the assumption that no such thing exists as a group in its own right. Yet this is not an epistemological claim but an ontological one and as Goldman points out, "if critics dispute the existence of such entities on ontological grounds, they will be picking a fight with a large community of philosophers" (Goldman 2010, p. 26). In a nutshell, this thesis has shown that, at a minimum, a departing from (A) will be required to a greater extent than both Goldman and Goldberg would be willing to admit individually. Alternatively, the process reliabilist framework that Goldman suggests (the "intra-personal" view), which clearly seeks to warrant as much of (A) as possible, asks for group accounts, like Rovane's, that question this and other features on even more radical grounds.

I hope that this thesis has proven rewarding in its own right in discussing this anti-individualistic strand and thereby presenting a more fine-grained discussion of the epistemic dimensions of List and Pettit's theory in process reliabilist terms. Moreover, if the objections made do indeed hold, it further shows that more concessions to epistemic anti-individualism are needed than both Goldman and Goldberg would each allow for.

References

- Alcoff, L., and Potter, E., eds. 1993. *Feminist Epistemologies*. New York, London: Routledge.
- Arrow, K. J. 1963. *Social Choice and Individual Values*. Yale University Press.
- Audi, R. 2006. "Testimony, Credulity, and Veracity." In *The Epistemology of Testimony*, edited by Jennifer Lackey and Ernest Sosa, 25–49. Oxford: Oxford University Press.
- Bird, A. 2010. "Social Knowing: The Social Sense of 'Scientific Knowledge.'" *Philosophical Perspectives* 24 (1): 23–56.
- . 2014. "When Is There a Group That Knows? Distributed Cognition, Scientific Knowledge, and the Social Epistemic Subject." In *Essays in Collective Epistemology*, edited by Jennifer Lackey, 42–63. Oxford: Oxford University Press.
- Bloor, D. 1991. *Knowledge and Social Imagery*. Chicago, London: University of Chicago Press.
- . 2005. "Toward a Sociology of Epistemic Things." *Perspectives on Science* 13 (3): 285–312.
- Bonjour, L. 1980. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5 (1): 53–73.
- BonJour, Laurence, and Sosa, E. 2003. *Epistemic Justification: Internalism Vs. Externalism, Foundations Vs. Virtues*. 450. Malden, MA ; Oxford: Blackwell Publishing.
- Booth, A. R. 2010. "The Theory of Epistemic Justification and the Theory of Knowledge: A Divorce." *Erkenntnis* 75 (1): 37–43.
- Bourdieu, P. 2004. *Science of Science and Reflexivity*. Chicago, London: University of Chicago Press.
- Bovens, L., and Rabinowicz, W. 2006. "Democratic Answers to Complex Questions: An Epistemic Perspective." *Synthese* 150 (1): 131–53.
- Bratman, M. 1999. *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge, New York: Cambridge University Press.
- Briggs, R. 2012. "The Normative Standing of Group Agents." *Episteme* 9 (03): 283–91.
- Burge, T. 1993. "Content Preservation." *The Philosophical Review* 102 (4): 457–88.
- Burke, P. 2000. *A Social History of Knowledge*. Cambridge: Polity Press.
- . 2013. *A Social History of Knowledge II: From the Encyclopaedia to Wikipedia: 2*. 1st ed. Cambridge: Polity Press.
- Cariani, Fabrizio. 2011. "Judgment Aggregation." *Philosophy Compass* 6 (1): 22–32.
- . 2012. "Epistemology in Group Agency; Six Objections in Search of the Truth." *Episteme* 9 (03): 255–69.
- Clark, A., and Chalmers, D. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19.
- Coady, C. A. J. 1992. *Testimony: A Philosophical Study*. Oxford: Oxford University Press.
- Corlett, J. A. 2007. "Analyzing Social Knowledge." *Social Epistemology* 21 (3): 231–2013.
- Dennett, D. C. 1971. "Intentional Systems." *The Journal of Philosophy* 68 (4): 87–106.
- . 1987. *The Intentional Stance*. Cambridge, MA: MIT Press.
- Dewitt, A. 2012. "Group Agency and Epistemic Dependency." *Episteme—a Journal of Individual and Social Epistemology* 9 (3): 235–44. doi:10.1017/epi.2012.13.
- Dietrich, F. 2008. "The Premises of Condorcet's Jury Theorem Are Not Simultaneously Justified." *Episteme* 5 (01): 56–73.

- Dietrich, Franz, and List, C. 2004. "A Model of Jury Decisions Where All Jurors Have the Same Evidence." *Synthese* 142 (2): 175–202.
- . 2006. "Arrow's Theorem in Judgment Aggregation." *Social Choice and Welfare* 29 (1): 19–33.
- . 2009. "The Impossibility of Unbiased Judgment Aggregation." *Theory and Decision* 68 (3): 281–99.
- . 2010. "The Aggregation of Propositional Attitudes: Towards a General Theory." *Oxford Studies in Epistemology* 3.
- . 2012. "Propositionwise Judgment Aggregation: The General Case." *Social Choice and Welfare* 40 (4): 1067–95.
- . 2013. "Propositionwise Judgment Aggregation: The General Case." *Social Choice and Welfare* 40 (4): 1067–95.
- Dunn, J. 2012. "Reliabilism: Holistic or Simple?" *Episteme* 9 (03): 225–33.
- Estlund, D. M. 1994. "Opinion Leaders, Independence, and Condorcet's Jury Theorem." *Theory and Decision* 36 (2): 131–62.
- Fagan, M. 2012. "Collective Scientific Knowledge." *Philosophy Compass* 7 (12): 821–31.
- Faulkner, P. 2000. "The Social Character of Testimonial Knowledge." *The Journal of Philosophy* 97 (11): 581–601.
- Fleck, L. 1980. *Entstehung und Entwicklung einer wissenschaftlichen Tatsache: Einführung in die Lehre vom Denkstil und Denkkollektiv*. Edited by Lothar Schäfer and Thomas Schnelle. 10th ed. Frankfurt am Main: Suhrkamp Verlag.
- Foley, R. 2004. "A Trial Separation Between the Theory of Knowledge and a Theory of Justified Belief." In *Sosa and His Critics*, edited by John Greco, 59–71. Oxford: Wiley-Blackwell.
- Foucault, M. 1972. *Archaeology of Knowledge*. London: Routledge.
- Fricker, E. 1995. "Testimony (Book Review) Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony." *Mind* 104: 393–411.
- Fricker, M. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford; New York: Oxford University Press.
- . 2012. "Group Testimony? The Making of A Collective Good Informant." *Philosophy and Phenomenological Research* 84 (2): 249–76.
- García-Carpintero, M., and Kölbel, M. 2008. *Relative Truth*. Oxford: Oxford University Press.
- Gaus, G. 2012. "Constructivist and Ecological Modeling of Group Rationality." *Episteme* 9 (03): 245–54.
- Gettier, E. L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23 (6): 121–23.
- Gilbert, M. 1987. "Modelling Collective Belief." *Synthese* 73 (1): 185–204.
- . 1990. "Walking Together: A Paradigmatic Social Phenomenon." *Midwest Studies in Philosophy* 15 (1): 1–14.
- . 1992. *On Social Facts*. Princeton: Princeton University Press.
- . 2004. "Collective Epistemology." *Episteme* 1 (2): 95–107.
- Goldberg, Sanford. 2007. *Anti-Individualism: Mind and Language, Knowledge and Justification*. Cambridge: Cambridge University Press.
- . 2008. "Reliabilism in Philosophy." *Philosophical Studies* 142 (1): 105–17.
- . 2010. *Relying on Others: An Essay in Epistemology*. Oxford: Oxford University Press.

- Goldman, A. I. 1979. "What Is Justified Belief?." In *Justification and Knowledge*, edited by George Pappas, 1–25. Boston: D. Reidel.
- Goldman, Alvin I. 1983. "Review." *The Philosophical Review* 92 (1): 81–88.
- . 1986. *Epistemology and Cognition*. Vol. 98. 3. Cambridge, MA; London: Harvard University Press.
- . 1999. *Knowledge in a Social World*. Vol. 36. 2. Oxford University Press.
- . 2002. *Pathways to Knowledge: Private and Public*. Oxford; New York: Oxford University Press.
- . 2004. "Group Knowledge Versus Group Rationality: Two Approaches to Social Epistemology." *Episteme* 1 (01): 11–22.
- . 2008. "Reliabilism in Philosophy." *Philosophical Studies* 142 (1): 105–17.. 1980. "The Internalist Conception of Justification." *Midwest Studies In Philosophy* 5 (1): 27–52.
- . 2010. "Why Social Epistemology Is Real Epistemology." In *Social Epistemology*, edited by Adrian Haddock, Alan Millar, and Duncan Pritchard, 1–29. Oxford University Press, Usa.
- . 2011. "Toward a Synthesis of Reliabilism and Evidentialism? Or: Evidentialism's Troubles, Reliabilism's Rescue Package." In *Evidentialism and Its Discontents*, edited by Trent Dougherty. Oxford: Oxford University Press.
- . 2012a. "Reliabilism." In *Reliabilism and Contemporary Epistemology*, by Alvin Goldman, 68–94. Oxford, New York: Oxford University Press.
- . 2012b. *Reliabilism and Contemporary Epistemology: Essays*. Oxford, New York: Oxford University Press.
- . 2014. "Social Process Reliabilism: Solving Justification Problems in Collective Epistemology." In *Essays in Collective Epistemology*, edited by Jennifer Lackey, 11–41. Oxford: Oxford University Press.
- Greco, J. 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- Greco, J, and Turri, J. 2013. "Virtue Epistemology." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2013. <http://plato.stanford.edu/archives/win2013/entries/epistemology-virtue/>.
- Hacking, I. 2001. *An Introduction to Probability and Inductive Logic*. Vol. 9. 4. Cambridge: Cambridge University Press.
- Hájek, A. 2008. "Arguments For—or Against—Probabilism?" *British Journal for the Philosophy of Science* 59 (4): 793–819.
- Hakli, R. 2007. "On the Possibility of Group Knowledge Without Belief." *Social Epistemology* 21 (3): 249–2013.
- Hakli, R., and Negri, S. 2011. "Reasoning About Collectively Accepted Group Beliefs." *Journal of Philosophical Logic* 40 (4): 531–55.
- Hales, S. D. 2011. *A Companion to Relativism*. Wiley-Blackwell.
- Hazlett, A. 2012. "Higher-Order Epistemic Attitudes and Intellectual Humility." *Episteme* 9 (03): 205–23.
- Heathcote, A. 2012. "Gettier and the Stopped Clock." *Analysis* 72 (2): 309–14.
- Hindriks, F. 2014. "How Autonomous Are Collective Agents? Corporate Rights and Normative Individualism." *Erkenntnis* 79 (9): 1565–85.

- Hobbes, T. 1996. *Hobbes: Leviathan: Revised Student Edition*. Edited by Richard Tuck. Student edition. Cambridge, New York: Cambridge University Press.
- Hume, D. 2008. "Of Miracles." In *An Enquiry Concerning Human Understanding*, by David Hume, edited by Peter Millican, 72–96. Oxford: Oxford University Press.
- Knight, J., and Johnson, J. 1994. "Aggregation and Deliberation: On the Possibility of Democratic Legitimacy." *Political Theory* 22 (2): 277–96.
- Kölbel, M. 2004. "Faultless Disagreement." *Proceedings of the Aristotelian Society* 104: 53–73.
- Kopec, M. 2012. "We Ought to Agree: A Consequence of Repairing Goldman's Group Scoring Rule." *Episteme* 9 (2): 101–14.
- Kornblith, H. 2001. *Epistemology: Internalism and Externalism*. Oxford: Blackwell Publishers.
- Kripke, S. A. 2011. "Nozick on Knowledge." In *Philosophical Troubles. Collected Papers Vol I*, edited by Saul A. Kripke, 162–224. Oxford: Oxford University Press.
- Kuhn, T. S. 2012. *The Structure of Scientific Revolutions*. Chicago, London: University of Chicago Press.
- Kusch, M. 2002. *Knowledge by Agreement: The Programme of Communitarian Epistemology*. Oxford: Oxford University Press.
- . 2014. "The Metaphysics and Politics of Corporate Personhood." *Erkenntnis* 79 (9): 1587–1600.
- Kusch, M., Pauer-Studer, H. and Schmid, H. B. 2014. "Introduction." *Erkenntnis* 79 (9): 1563–1563.
- Lackey, J. 2008. *Learning From Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.
- . 2010. "Testimony: Acquiring Knowledge From Others." In *Social Epistemology: Essential Readings*, edited by Alvin I. Goldman and Dennis Whitcomb. Oxford: Oxford University Press.
- . 2014a. "A Deflationary Account of Group Testimony." In *Essays in Collective Epistemology*, edited by Jennifer Lackey, 64–96. Oxford: Oxford University Press.
- , ed. 2014b. *Essays in Collective Epistemology*. Oxford ; New York: Oxford University Press.
- Lahroodi, Reza. 2007. "Collective Epistemic Virtues." *Social Epistemology* 21 (3): 281–2013.
- List, C. 2003. "A Possibility Theorem on Aggregation over Multiple Interconnected Propositions." *Mathematical Social Sciences* 45 (1): 1–13.
- . 2004a. "A Model of Path-Dependence in Decisions over Multiple Propositions." *American Political Science Review* 98 (3): 495–513.
- . 2004b. "The Impossibility of a Paretian Republican? Some Comments on Pettit and Sen." *Economics and Philosophy* 20 (1): 65–87.
- . 2005. "Group Knowledge and Group Rationality: A Judgment Aggregation Perspective." *Episteme* 2 (01): 25–38.
- . 2006. "The Discursive Dilemma and Public Reason." *Ethics* 116 (2): 362–402.
- . 2011. "The Theory of Judgment Aggregation: An Introductory Review." *Synthese* 187 (1): 179–207.
- . 2014. "Three Kinds of Collective Attitudes." *Erkenntnis* 79 (9): 1601–22.
- List, C., and Pettit, P. 2002. "Aggregating Sets of Judgments: An Impossibility Result." *Economics and Philosophy* 18 (1): 89–110.
- . 2004. "Aggregating Sets of Judgments: Two Impossibility Results Compared." *Synthese* 140 (1-2): 207–35.

- . 2006. "Group Agency and Supervenience." *The Southern Journal of Philosophy* 44 (S1): 85–105.
- . 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.
- Longino, H. 2013. "The Social Dimensions of Scientific Knowledge." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2013. <http://plato.stanford.edu/archives-spr2013/entries/scientific-knowledge-social/>.
- Mathiesen, K. 2006. "The Epistemic Features of Group Belief." *Episteme* 2 (3): 161–75.
- . 2007. "Introduction to Special Issue of Social Epistemology on 'Collective Knowledge and Collective Knowers.'" *Social Epistemology* 21 (3): 209–213.
- Mcmyler, B. 2011. "Knowing Second Hand." In *Testimony, Trust, and Authority*, 45–76. Oxford: Oxford University Press.
- Mongin, P. 2008. "Factoring out the Impossibility of Logical Aggregation." *Journal of Economic Theory* 141 (1): 100–113.
- . 2012. "The Doctrinal Paradox, the Discursive Dilemma, and Logical Aggregation Theory." *Theory and Decision* 73 (3): 315–55.
- Moran, R. 2005. "Getting Told and Being Believed." *Philosophers' Imprint* 5 (5): 1–29.
- Nehring, K. 2003. "Arrow's Theorem as a Corollary." *Economics Letters* 80 (3): 379–82.
- Nozick, R. 1981. *Philosophical Explanations*. Harvard University Press.
- Palermos, S. O., and Pritchard, D. 2013. "Extended Knowledge and Social Epistemology." *Social Epistemology Review and Reply Collective*, no. 8: 105–20.
- Pauer-Studer, H. 2014. "A Constitutive Account of Group Agency." *Erkenntnis* 79 (9): 1623–39.
- Pettit, P. 1993. *The Common Mind*. Oxford: Oxford University Press.
- . 2006. "Joint Actions and Group Agents." *Philosophy of the Social Sciences* 36 (1): 18–39.
- . 2007. "Joining the Dots." In *Common Minds: Themes From the Philosophy of Philip Pettit*, edited by Geoffrey Brennan, Robert E. Goodin, and Michael A. Smith. Oxford: Oxford University Press.
- . 2010. "Groups with Minds of Their Own." In *Social Epistemology: Essential Readings*, edited by Alvin I. Goldman and Dennis Whitcomb, 242–68. Oxford: Oxford University Press.
- . 2014. "Group Agents Are Not Expressive, Pragmatic or Theoretical Fictions." *Erkenntnis* 79 (9): 1641–62.
- Pettit, P., and Schweikard, P. 2006. "Joint Actions and Group Agents." *Philosophy of the Social Sciences* 36 (1): 18–39.
- Pigozzi, G. 2006. "Belief Merging and the Discursive Dilemma: An Argument-Based Account to Paradoxes of Judgment Aggregation." *Synthese* 152 (2): 285–98.
- Reid, T. 2011. *An Inquiry into the Human Mind, on the Principles of Common Sense*. Reprint. Cambridge: Cambridge University Press.
- Rousseau, J.-J. 2011. *Rousseau: The Social Contract and Other Later Political Writings*. Edited by Victor Gourevitch. New. Cambridge, New York: Cambridge University Press.
- Rovane, C. 1998. *Bounds of Agency: An Essay in Revisionary Metaphysics: Essay in Revisionary Metaphysics*. Princeton, N.J.: Princeton University Press.
- . 2011. "Relativism Requires Alternatives, Not Disagreement of Relative Truths." In *A Companion to Relativism*, edited by Steven Hales, 31–52. Chichester: Wiley-Blackwell.
- . 2014. "Group Agency and Individualism." *Erkenntnis* 79 (9): 1663–84.

- Runciman, D. 2008. *Pluralism and Personality of State*. Revised. Cambridge, UK; New York: Cambridge University Press.
- Schmid, H. B. 2008. *Plural Action. Essays in Philosophy and Social Science*. Dordrecht: Springer.
- . 2012. *Wir-Intentionalität: Kritik des ontologischen Individualismus und Rekonstruktion der Gemeinschaft*. 1st ed. Freiburg im Breisgau: Verlag Karl Alber.
- . 2014. "Expressing Group Attitudes: On First Person Plural Authority." *Erkenntnis* 79 (9): 1685–1701.
- Schmid, H. B., Sirtes, D. and Weber, M. 2011. *Collective Epistemology*. Berlin, Boston: Walter de Gruyter.
- Schmitt, F. F. 1994. *Socializing Epistemology: The Social Dimensions of Knowledge*. Lanham, Boulder, New York: Rowman and Littlefield.
- Searle, J. R. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Vol. 20. 79. Cambridge: Cambridge University Press.
- . 1995. *The Construction of Social Reality*. London: Penguin Books.
- . 2002. "Collective Intentions and Actions." In *Consciousness and Language*, by John R. Searle, 90–105. Cambridge ; New York: Cambridge University Press.
- Sen, A. K. 1966. "A Possibility Theorem on Majority Decisions." *Econometrica* 34 (2): 491–99.
- Solomon, M. 2006. "Groupthink versus The Wisdom of Crowds: The Social Epistemology of Deliberation and Dissent." *The Southern Journal of Philosophy* 44 (S1): 28–42.
- Sosa, E. 2006. "Knowledge: Instrumental and Testimonial." In *The Epistemology of Testimony*, edited by Jennifer Lackey and Ernest Sosa, 116–23. Oxford: Oxford University Press.
- . 2009. *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford: Oxford University Press.
- Stalnaker, R. 1996. "Varieties of Supervenience." *Noûs* 30 (January): 221–41.
- Sylvan, K. L. 2012. "How to Be a Redundant Realist." *Episteme* 9 (3): 271–82.
- Thagard, P. 1997. "Collaborative Knowledge." *Noûs* 31 (2): 242–61.
- Tollefsen, D. 2004. "Collective Epistemic Agency." *Southwest Philosophy Review* 20 (1): 55–66.
- . 2006. "Group Deliberation, Social Cohesion, and Scientific Teamwork: Is There Room for Dissent?" *Episteme* 3 (1-2): 37–51.
- . 2007. "Group Testimony." *Social Epistemology* 21 (3): 299–2013.
- Tuomela, R. 2004. "Group Knowledge Analyzed." *Episteme* 1 (2): 109–27.
- . 2010. *The Philosophy of Sociality: The Shared Point of View*. Oxford, New York: Oxford University Press.
- Wray, K. B. 2001. "Collective Belief and Acceptance." *Synthese* 129 (3): 319–33.
- . 2010. "Introduction: Collective Knowledge and Science." *Episteme* 7 (03): 181–84.
- Wright, C. 2006. "Intuitionism, Realism, Relativism and Rhubarb." In *Truth and Realism*, edited by Patrick Greenough and Michael Lynch, 38–60. Oxford: Clarendon Press.

Appendix

List of Tables:

Table 1 Discursive Dilemma.....	51
Table 2. Positive and Negative Truth-Tracking and Indicating Reliabilities	62

List of Figures:

Figure 1 Relationship between Members and a Group.....	36
Figure 2 Aggregation Functions	53

Zusammenfassung

Es existieren viele Weisen die kollektive Dimension von Wissen zu thematisieren. Diese Arbeit setzt sich mit epistemischen Implikationen einer spezifischen These des Zusammenhangs zwischen Wissen (oder gerechtfertigter Überzeugung) und Kollektivität auseinander. Im Zentrum steht dabei die Annahme, dass es gerechtfertigte Gruppenüberzeugungen geben kann. Gruppen werden dabei als eigenständige Entitäten betrachtet die nicht auf ihre Mitglieder reduzierbar sind.

Im ersten Teil der Arbeit wird nicht nur diese ontologische Annahme vorgestellt, sondern auch der epistemologische Rahmen zu dem sich diese Arbeit in ihrer Auseinandersetzung mit gerechtfertigter Gruppenüberzeugung verpflichtet abgesteckt. In ontologischer Hinsicht wird dabei Non-Summativismus [*non-summativism*] bezüglich Gruppen angenommen, in epistemologischer Hinsicht wird festgehalten, dass sich diese Arbeit auf erkenntnistheoretischen Externalismus, und insbesondere auf Verlässlichkeitstheorien die auf Prozesse fokussieren [*process reliabilism*] beschränkt.

Im zweiten Teil dieser Arbeit wird eine spezifische non-summativistische Gruppentheorie vorgestellt und kritisch diskutiert, nämlich Christian List und Philip Pettits Theorie von Gruppenakteurschaft [*group agency*]. Neben einer generellen Einführung in ihr Modell, wird der Fokus vor allem auf die epistemologischen Implikationen dieser Theorie gelegt. Ihr Vorschlag eines auf Wahrheit abzielenden [*truth-tracking*] Modells weist einige potentiell problematische Aspekte auf, die aufgezeigt und diskutiert werden.

In einem dritten und letzten Teil wird ein erkenntnistheoretischer Ansatz vorgestellt – Alvin Goldmans Theorie eines „social process reliabilism“ – der es ermöglichen könnte, die identifizierten Probleme zu vermeiden. Während Goldman davon ausgeht, dass sein Ansatz mit List und Pettits Theorie von Gruppenakteurschaft kompatibel sei, wird in einem Ausblickskapitel gezeigt, dass diese Annahme nicht haltbar ist. Zwei Optionen um diesem Umstand zu begegnen werden vorgestellt. Beide dieser Optionen haben zur Folge, dass diverse individualistische Annahmen die Goldmans Theorie charakterisieren in Frage gestellt werden müssen.

Abstract

There exist many ways in which it would be sensible to analyse the social dimension of belief formation. This thesis discusses the epistemic implications of one specific way in which a belief can be seen as something collective. It concerns itself with the justification of group beliefs, that is to say, with beliefs that are held by a group.

The first part of this thesis not only introduces a general framework within which group beliefs are discussed, but also restricts itself to a particular epistemological framework. For social ontology, non-summativism about groups is assumed. According to non-summativism, groups exist as entities in their own right. For epistemology, this thesis commits itself to epistemic externalism and, more precisely, to process reliabilism.

The second part of this thesis discusses one particular account of non-summativism, namely Christian List and Philip Pettit's model of group agents. After a fairly descriptive introduction to their account, their model of truth tracking in groups is presented and critically examined. A closer discussion reveals several potential problems with List and Pettit's tracking account.

The third and final part of this thesis attempts to present a solution to these problems by introducing the process reliabilist framework for analysing justified group belief recently suggested by Alvin Goldman (2014). In an overview chapter, some potential problems with Goldman's social process reliabilist account are presented and two potential solutions are suggested. As it turns out, each of these solutions entails a deviation from individualism to a greater degree than social process reliabilism suggests.

Lebenslauf

Katharina Bernhard, BA

Ausbildung:

2012– aktuell: MA Philosophie, Universität Wien.

2012: BA Philosophie, Universität Wien

2003–2013: Studium der Philosophie und zeitweises Studium der Internationalen Entwicklung (2003–2007), Germanistik (2004), Vergleichende Literaturwissenschaft (2007–2013)

2003: Matura, BRG 7, Bundesrealgymnasium Kandlgasse, Wien

1995–2003: AHS, BRG 7, Bundesrealgymnasium Kandlgasse, Wien

Akademische Anstellungen:

SoSe 2012 – WS 2013: Studienassistentin für H.B. Schmid, Universität Wien

SoSe 2013: Zusätzliche Tutoriatsstelle für StEOP, H.B. Schmid, Universität Wien

SoSe 2014: Tutoriatsstelle für H.B. Schmid (als Tutorin für StEOP), Universität Wien

WS 2014: Tutoriatsstelle für H.B. Schmid (VO Praktische Rationalität), Universität Wien

Mitgliedschaften:

2012–aktuell: Mitglied des *Wiener Forums for Analytische Philosophie* (WFAP)

Außeruniversitäres:

2003–aktuell: arbeitstätig bei Bernhard-AV GmbH