

COLLECTIVE INTENTIONALITY AND SOCIAL AGENTS

Raimo Tuomela

University of Helsinki

Helsinki, Finland

I INTRODUCTION

I.1

In this paper I will discuss a certain philosophical and conceptual program -- that I have called philosophy of social action writ large -- and also show in detail how parts of the program have been, and are presently being carried out. In current philosophical research the philosophy of social action can be understood in a broad sense to encompass such central research topics as action occurring in a social context (this includes multi-agent action); shared we-attitudes (such as we-intention, mutual belief) and other social attitudes expressing collective intentionality and needed for the explication and explanation of social action; social macro notions, such as actions performed by social groups and properties of social groups such as their goals and beliefs; social practices, and institutions (see e.g. Tuomela, 1995, 2000a, 2001). The theory of social action understood analogously in a broad sense would then involve not only philosophical but all other relevant theorizing about social action.¹

I will below will sketch how to build the macro aspects of the social world conceptually out of what I will call the “*jointness*” level with special emphasis on notions expressing *collective intentionality* and show how human beings have constructed the central parts of the social world. If we follow sociologists in speaking of macro-, meso-, and micro-level concepts in the social sciences, then it can be said that I will try to argue that the macro level can ontically be constructed out of the meso and micro levels, given a suitably rich conceptual framework, the so-called “framework of agency”, including not only intentional single-agent notions but such jointness level notions as shared we-attitudes, joint and collective action, various notions of social interaction and dependence, as well as normative notions (e.g. social norm, obligation, commitment). Speaking in terms commonly used in Artificial Intelligence, I will argue that essentially the belief-desire-intention (BDI) conceptual and ontic architecture formulated for the jointness level and fortified with the mentioned kinds of normative notions will accomplish the task or at least the central part of the task. My arguments below will not suffice for all of what I just claimed, but I will show in detail how e.g. group properties can be analyzed in terms of jointness level notions and make some

sketchy remarks concerning some remaining problems. I have chosen to emphasize in particular that the so-called we-mode/I-mode distinction as applied to attitudes and actions is important and that the macro-level social notions essentially will depend on the use of we-mode notions.

While this is to large extent a programmatic paper, I should immediately add that much of the program already has been carried out (see Tuomela, 1995, 2000a, 2001, Tuomela and Balzer, 1999, Balzer and Tuomela, 1997, 2001). At various stages of this work slightly different terminology has been used, but I will below try to show the connections.

My detailed treatments will be interesting also in themselves over and above their role in arguing for some grand thesis or other. Before starting, I must regretfully say that such important topics as the nature and import of social practices, social institutions, and the development of these notions cannot be properly handled below, if at all (but see Tuomela, 2001). I will not in fact even discuss very much the phenomena of collective intentionality – viz. phenomena such as joint intentions and mutual beliefs. There is a section on shared we-attitudes, which form my most central category of states expressing collective intentionality, but mainly due to lack of space I will not go deep into the discussion of particular topics – such as joint intentions and mutual beliefs (for which see e.g. Tuomela, 1995, Chapters 1 and 3, 2000a, 2000c, 2001). Another thing that I cannot properly discuss below are the criteria of adequacy of my analyses. Most importantly, the central criteria will not only be “hermeneutic” but also functional – it must be shown how the analyses correspond to our common sense notions and how they make the analyzed phenomena work under their new guises. As said, the main sections of the paper will concentrate on the centrality of notions expressing we-mode collective intentionality and it will show their role in building social macro notions such as group attitudes.

I.2.

To end this introductory section, I will present a number of philosophical theses about the nature of the social world that I take to be right and that I have at least to some extent defended elsewhere. Nevertheless, there will also be new points to be made that I have not previously said in print. My new points here will mainly be related to the importance of collective intentionality as “shared we-attitudes” and “collective acceptance”, as well as to the centrality of the “we-mode/I-mode” distinctions. What connects the theses is that they involve the so-called "Collective Acceptance" view of sociality. According to this view certain central parts of the social world, such as social institutions, are man-made and based on collective acceptance. Here collective acceptance amounts to coming to hold and holding a relevant shared “we-attitude” (basically we-intention and/or we-belief). (The details of the model can be found in Tuomela and Balzer, 1999, and Tuomela, 2001.)

Here are the theses:

T1) Notions of collective intentionality are central and needed for understanding society and social life in general. Collective intentionality most centrally includes shared we-attitudes, which include e.g. joint intentions, wants, commitments, and we-beliefs.

T2) The so-called Collective Acceptance account makes use of shared we-attitudes -- especially shared we-intentions and we-beliefs -- and accounts for social norms and social institutions (their emergence, maintenance, and renewal) and various other collective social phenomena. In effect, collective acceptance amounts to coming to hold and holding relevant shared we-attitudes. Coming to collectively accept an item involves conation, viz. collective intention. However, collective acceptance need not be planned but can be “acceptance-in-action”.

T3) In general, it is important to distinguish between attitudes and actions in the we-mode and those in the I-mode. The former are essentially the kinds of positional attitudes and actions that a member has when having the attitude or, respectively acting, “*qua* a group member” in the fullest sense, whereas I-mode attitudes do not thus depend on group membership.² Accordingly, one can speak of we-mode and I-mode social agency. An important application is cooperation, for it is central to distinguish between we-mode cooperation and I-mode cooperation (cf. the resulting two-faced view of cooperation created in Tuomela, 2000a).

T4) It is important to distinguish also between normative, group-binding and non-normative group properties (attitudes and actions). Thus there are normative macro-social properties such as group beliefs, group goals, and institutional social practices which are to be distinguished from non-normative group beliefs, group goals, and group practices (cf. Tuomela, 1995).

T5) The Collective Acceptance account distinguishes between normative and non-normative group properties (attitudes and actions) by taking the former, in contrast to the latter, to be based on the group’s “authority system”, viz. its joint-will-formation system, which has been authorized by the group, typically in the sense of authorizing some “operative” members for joint-will formation. An exercise of the group’s authority system in this context requires the obtaining of special socially and normatively right circumstances, including “positional” constraints, with the result that the operative members function properly *qua* members of the group. The non-operative agents of the group have a *prima facie* obligation to “tacitly accept” (although only in a weak sense of acceptance) what the operative members here do.

In the special case of creating and upholding group-binding attitudes and actions collective acceptance requires explicit or implicit agreement making concerning the item, say proposition p, in question and results in collective commitment concerning p. Here p can express a view (e.g. “the

earth is flat”), which gives a group belief, or it can express a plan (e.g. “we will have a swim together in the afternoon” or “we will overthrow capitalism”), which gives group intention.

T6) In attempts to understand social life it is, furthermore, central to make a distinction between the (metaphorically understood) "agreement-aspects" and "belief-aspects" of both collective acceptance and the social realm. The former are in the last analysis based on the notion of *agreement making* and a public or objective *we-intention* while the latter consist of the vaguer and "softer" notions that mere *mutual belief* (and action based on mere mutual belief) can serve to create and sustain. Obviously, “rule norms” (and rule-based tasks) are basic building blocks of the former kind of aspect of the social world, while norms based on mutual expectations (as well as expectation-based social roles), for instance, are central elements of the latter kind of aspect.

T7) Both institutional and non-institutional *social practices* involve collective intentionality in the sense of being based on shared we-attitudes in the sense of we-attitude contents serving as partial reasons for such activities. Such reasons need not be reflected on or even focused on when acting and they accordingly need play only the guiding role of a presupposition.

T8) As a special case, the Collective Acceptance account serves to characterize the central features of social institutions, which are basically reflexive social practices in a normative context. A social organization such as a business company can be mentioned as a special, developed case of a social institution that also falls within the scope of the Collective Acceptance account.

T9) The Collective Acceptance account also serves diachronically to account for the emergence, maintenance, and development of social practices. Basically, collective attitudes under suitable circumstances give rise to collective social action and social practices purported to satisfy those attitudes. Some of those practices become norm-governed and institutionalized either instantaneously or in the course of time due to performative collective acceptance. The results and consequences of such actions feed back into the collective attitudes and the institutional background (once it has been created), and thus maintenance and renewal of social practices and institutions will result. Unintended and unforeseen consequences also have a function in the dynamic Collective Acceptance model, as also they are fed back into the motivational and structural background.

T10) Group attitudes and actions, and indeed all macro-level social notions (including, e.g. such social notions as cohesion, status, and stratum) *ontically* supervene (in a precise, technical sense) on lower level, viz. “meso-level, notions. (However, group *concepts* can be argued not to be reducible to lower level concepts, which makes conceptual holism about social notions a tenable view.)

T11) The Collective Acceptance view accounts for all man-made (thus “constructivist”) aspects of social life, including social values and morality.

I will discuss the contents of many the theses in the text to follow, but without mentioning the theses explicitly. All of the theses have been discussed at least to some extent in my previous work.

II NORMATIVE GROUP BELIEFS AND OTHER NORMATIVE GROUP ATTITUDES

Common sense attributes goals and beliefs to groups, and I will respect this idea here. Goals and beliefs are thus ascribable (somewhat metaphorically) to groups, both formal and informal, structured and unstructured. Drawing on the developments in Tuomela, 1995, I will now discuss group beliefs in detail. As will be seen, the “positional” account to be given of them and other group attitudes and of group actions is ultimately based on relevant collective acceptance, and thus the model below is a special case of the Collective Acceptance account developed in Tuomela and Balzer, 1999, and Tuomela, 2001.

Concentrating on the group belief case, consider the following examples:

- a) The Government believes that war against Ruritania will begin soon.
- b) The Catholic Church believes that miracles happen.
- c) Texaco believes that children should be seen, not hurt.
- d) The Communist Party of Ruritania believes that capitalist countries will soon perish (but none of its members really believes so).
- e) Our family believes that the schools in this country are inefficient.
- f) The team believes that it will win today's game.
- g) The Finns believe that sauna originated in Finland.
- h) This mob believes that Smith is a traitor.
- i) Europeans believe that face-to-face discussants should keep at least half a meter apart from each other.

I will below develop an account of group beliefs which will accommodate cases a)-i) and which, indeed, is meant to cover all possible cases of group belief. My main account will concern a group's beliefs, viz., cases in which a group as a whole can be taken to believe something, say p. These will be called (proper) group beliefs or normatively group-binding group beliefs. There is also another kind of social belief that can be called a group belief. This is belief characterizable as a

shared we-belief; also social collectives which are not groups capable of action can, of course, have shared we-beliefs.

Two features that I have emphasized in my developments (in Tuomela, 1995, Chapters 5-7, especially) are the topic-relative distinction between operative and non-operative members of a social collective and the notion of correct social and normative circumstances. As argued there, these are also highly relevant to the problem of ascribing beliefs and views to groups and to the problem to what extent groups can even metaphorically be taken to resemble persons. The notions of goal, belief, and action are linked in the case of a group to approximately the same degree as in the individual case. In the latter case their interconnection is well established; given that the person-analogy applies to groups (which we will accept here in the case of these basic notions), these notions apply to groups as well. More broadly and loosely, a group's having a goal and a relevant means-belief related to it is or involves a disposition to act in the pursuit of the goal; and if action requires the mentioned notion of a socially and normatively right circumstance and the notion of an operative member, then so does a disposition to act, otherwise group action manifesting it could not -- on conceptual grounds -- come about or even get initiated. I will clarify these notions below.

The operative members in the cases of group actions, group goals, and group beliefs are those actors, goal-formers, and belief-formers by virtue of whom, respectively, actions, goals and beliefs are attributed to groups. Operative members can include non-member "representatives" such as attorneys, delegates, and hired persons. The collective action of operative members (or the action of the operative member, if there is only one of them) in this sense is sufficient but perhaps not necessary for the group action or attitude to come about.

Typically the operative members of a group g can be regarded as -- or at least will include -- the *leaders* of g . While the notion of an operative member and that of an (implicitly or explicitly) authorized member are different, in the case of our central concern -- viz. in the case where the operative members are "internally" authorized by the group members -- the set of operative members will coincide with, or at least belong to, the set of authorized members. Accordingly, the operative agents of g can be said to be authorized, institutionally or by informal or only by implicit agreement, to act on behalf of the group. Authorization here means commission authority, possibly in an implicit sense only (see Tuomela, 1995, Chapter 4). Note that we are dealing with two central distinctions here. The first is the operative/non-operative distinction, where we in principle can have among the operatives of g both members and non-members of g while the non-operatives are always members of g . The second distinction is between the case where the operatives get their authority to act from the members of g , the case where the authority comes from outside the group, so to speak (e.g., group-external laws which are in force in g), and the case where the authority is

created by other means (cf. dictatorship). My phrasing below will, for simplicity's sake, focus on the type of case where the authority of the operative members comes from the group members and where, furthermore, all the operative members are group members.

As an example, consider a state's making a pact with another state. This takes place by, say, the Cabinet ministers' jointly accepting the pact and the Prime Minister's signing it. Most citizens of the state do nothing relevant here, we may assume; the Cabinet represents them. But the state would hardly act fully intentionally as a whole unless the (non-operative) citizens were aware and at least "tacitly" or "silently" accepted the fact that there are operative agents acting on their behalf (generally and also in that specific way), although they may not know who they are and in what tasks they act nor exactly what relevant act they are performing (imagine a secret pact made in the name of the state by the Cabinet members without anybody else's faintest knowledge of it). In some cases it may be appropriate to require that the non-operative members have a relevant conditional group-intention, where the condition might be that they are elected, or otherwise become, operative members, thus changing status from being non-operative ones.

In the case of intentional actions performed by collectives also the non-operative agents, so to speak, "passively" participate in virtue of tacitly accepting (or at least being obligated so to accept) what the operative members do for the group (see below and Tuomela, 1995, Chapter 6 for more details).

The right social and normative circumstances are those in which the members of collectives act and form views in their positions (and in their right tasks within them). Positions can be regarded basically as collections of tasks, based on social rule norms, and social roles, based on "proper", expectation-based social norms (see Tuomela, 1995, Chapters 1,5, and 8 for these notions). The rules defining the tasks in question can be either formal (resembling laws and statutes) or informal (based on informal group-agreement). I will below speak as if the collectives and positions in them were characterized at least in part in terms of the formal or informal rules which are in force in them. (However, the present approach can handle group beliefs also in the cases of a group without previous rules, as long as there is at least an implicit "authority-incorporating" system, viz. "authority system", for creating joint decisions and commitments in the group.)

Accordingly, as argued in detail in Chapters 5-7 of Tuomela (1995), the notions of the socially and normatively right circumstances and the notion of an operative member are crucial in accounting for the functioning of formal and organized collectives such as corporations. In their case the statutes, by-laws and other relevant rules of the collective can be shown to connect goals (interests, purposes, and whatever subtypes of goals are at stake), beliefs (or views), and actions. To have successful group action we must also require that the right circumstances (as specified or

presupposed in the rules) obtain; and group beliefs about those circumstances typically mediate between goals and actions, in analogy with the single-agent case. Indeed, in the case of typical formal collectives, certain position-holders are required by the constitutive rules of the collective to set goals and accept views for the collective. So the main argument here is that because formal collectives like corporations require the previously mentioned technical notions for an adequate account not only of their actions but also their goals and beliefs, we need these technical notions, suitably liberalized, also in other, less formal cases. (In the case of mobs, for instance, all the members may be operative members, and we are dealing with a special case.)

In our above example of the treaty proposal, the Prime Minister's signing does not qualify if the content of the treaty is unlawful, or if it is signed in circumstances not qualified by law. In the context of group action he can intentionally do everything which he has the authority to do and which, accordingly, does not violate the rules defining the position and the contextually relevant tasks of a Prime Minister. In general, we should demand (at least in the case of structured groups with rules) that the circumstances *C* be such that the following at least "prototypically" holds: the ought-to-be rules (saying, e.g., what the administration ought be like), ought-to-do rules (specifying what position-holders ought to do in various circumstances), and may-do rules (specifying analogously which actions are permitted) characterizing an organization may not at least intentionally be violated by the position-holders. But even more must be required (cf. Tuomela, 1995, Chapter 5). If we concentrate on a single operative member A_i we may say the following. Assume that this agent is a position-holder in the collective. His actions as a position-holder are governed by a) formal or informal rules specifically concerning his tasks and rights (and/or actions fulfilling them), b) general constitutive rules concerning the purposes and functions of the collective which are not directly concerned with positional tasks and rights but which underlie them, and c) proper social norms (typically covert, unwritten norms) specifying his social roles. What is central for the purposes of this chapter is task-fulfilling behavior -- in contrast to actions performed for fulfilling the demands of social roles. The positional tasks and rights that we will be concerned with are based on (formal or informal) social rules in force in the collective. (These rules are basically rules of kind a) but rules b) cannot be violated either, and when speaking of rules specifying tasks and of the resulting rule-tasks below we will assume thus.) To a position (a "p-category") in a collective, accordingly, will be attached rule-tasks, viz., states to be brought about or maintained, as required by social rules of the aforementioned kind. It also contains rights, although, for simplicity's sake, we will in our formulations often suppress mention of rights. (Note that even if role-tasks will in general be attached to a position they are here not relevant -- see Tuomela, 1995, Chapter 5.)

One of the key concepts of our analysis of group beliefs (and goals) is the notion of accepting a proposition (or, if you prefer, a sentence). Accepting here involves agreeing to a proposition (or sentence), say *p*, and involves taking *p* to be true. As contrasted with the belief that *p*, accepting can be performed as an action, something one can do at will or on purpose, while what can be called "experiential" believing is a non-actional, dispositional mental state that in typical cases is brought about causally (cf. perceptual beliefs). Acceptance need not always be intentional, however, but this does not annihilate the fact that accepting and believing (taken in this experiential sense) are different notions. However, the acceptance and the pure belief aspects in actual mental states may not be more than analytically distinguishable. Thus acceptance beliefs can in addition have experiential features (related to a feeling that the believed state of affairs is real and, in strong cases, to being convinced about it). The acceptance aspect of beliefs again relates to the rational and cognitive features of beliefs, e.g., to their use as premises in reasoning and argumentation and to their logical properties -- such as logical closure under conjunction, with some qualifications perhaps. (See Cohen, 1992, Tuomela, 1992, 1995, and 2000b for discussion.)

I will below speak of experiential belief and belief in the sense of acceptance as true, both regarded as analytical ideal types. (Note that experiential belief also is truth-related.) I will argue that group beliefs and what will be called "positional beliefs" (viz., we-mode beliefs) may involve acceptance in a wide sense not necessarily related to truth and thus may not involve acceptance of propositions as true nor (truth-related) experiential belief. The notion of accepting a proposition as true will be called "acceptance in the narrow sense." Actually the distinction between (narrow) acceptance belief and experiential belief is not central below, in contrast to the distinction between ordinary truth-related belief versus "belief" not so related, viz., wide acceptance. Joint acceptance in this context will involve making a joint agreement.

Let me now reproduce the *positional account of group-binding group beliefs* (or positional joint agreement-account of group beliefs) as it has been formulated in Tuomela, 1992 and 1995, Chapter 7. The model relies on some technical notions which cannot be fully explicated here. Thus not only is the operative/non-operative distinction assumed but also the notions of a position in the group and the tasks (rule-tasks and roles tasks) related to positions, and the notion of correct social and normative circumstances is relied on. To indicate what is at stake, the Prime Minister of a state occupies a position in it with which various ought-to-be, ought-to-do, may-be, and may-do norms are attached. These norms serve to define the position-holder's (here Prime Minister's) tasks and indeed the content of his position and its normative relationships to other positions in the collective. When the Prime Minister acts in his position he acts *qua* a group member in this position, and when functioning correctly he has to respect the limits set to his functioning as a position-holder. When

he, for instance, signs a contract for the state, or accepts a view or a goal for it, he must act in the socially and normatively right circumstances (e.g. in a cabinet meeting or whatever is on the books) and not e.g. act in a movie representing himself.

The positional account of group beliefs can be sketchily summarized as follows. A group is taken to believe something p if it accepts p as its view. This can only be the case if the group members or some of them, the operative ones, collectively or jointly accept p for the group. When they so do they must be acting correctly *qua* group members, viz. functioning in their positions in the group when the right social and normative circumstances obtain. The non-operative members of the group must tacitly accept, or at least put up with, what the operative members accept as the group' views. They need not even have detailed knowledge about what is so being accepted. A more detailed account is the following:

(BC) *Group g believes that p in the normative, group-binding sense in the social and normative circumstances C if and only if in C there are operative members A_1, \dots, A_m of g in respective positions P_1, \dots, P_m such that*

1) the agents A_1, \dots, A_m , when they are performing their social tasks in their positions P_1, \dots, P_m and due to their exercising the relevant authority system ("joint intention formation" system) of g , (intentionally) jointly accept p and because of this exercise of the authority system they ought to continue to accept and positionally believe it;

2) there is a mutual belief among the operative members A_1, \dots, A_m to the effect that 1);

3) because of 1), the (full-fledged and adequately informed) non-operative members of g tend to tacitly accept - or at least ought to accept - p , as members of g ;

4) there is a mutual belief in g to the effect that 3).

Social tasks here will have to include rule-tasks. Put in general terms, the notion of joint acceptance involves the joining of the participants' wills into a group will, and thus a joint commitment. Accordingly, acceptance must in clause 1) be taken to involve agreement-making. Moreover, agreement must always be intentional. Thus we are dealing with intentional acceptance, although such acceptance may be less than (at least fully) conscious. (The concept of tacit acceptance of clause 3) is a much weaker notion; cf. Tuomela, 1995, Chapter 6.) The notion of joint acceptance of something as g 's view in the sense of clause 1) can be spelled out in terms of joint positional beliefs:

(*) Agents A_1, \dots, A_m *jointly accept p as the view of g* (or *jointly accept p for g to express its view*) in the social and normative circumstances C , if and only if, in C these agents (in their respective positions) jointly accept p because of their relevant exercise of the authority system of g .

We are here dealing with positional acceptance-belief - a mental state resulting from wide acceptance; and this is what accepting a belief means in (*). Joint acceptance here then comes to mean joint or shared acceptance-belief. Here joint acceptance must be an intentional proper joint action, and thus a mutual belief (not necessarily conscious) in the joint acceptance will exist. An acceptance-belief in this positional sense need not be a belief in the ordinary truth-related sense, but mutual belief can be regarded as a truth-related belief.

Of our earlier specific example sentences, my present analysis of positional normative group beliefs will cover a), b), c), d) (as an extreme case), e), and f) (in some cases also h).

Closely similar analyses can be given of group goals and intentions and even group actions. As an illustration, here is my account of action intentions attributed to a group (Tuomela, 1995, p. 286):

(GI) A group *g* *intends* to perform or bring about X in the social and normative circumstances C if and only if in C there are operative agents A_1, \dots, A_m such that

1) the operative agents A_1, \dots, A_m , when performing their social tasks in their positions P_1, \dots, P_m and due to their exercising the relevant authority system of *g*, have come intentionally jointly to accept "We will do X" and thus to we-intend (or at least to be disposed to we-intend) to perform X (and hence are committed to endorsing this intention), and, furthermore, because of this exercise of the authority system they also ought to continue to accept this we-intention;

2) because of 1), the (full-fledged and adequately informed) non-operative members of *g* tend to tacitly accept the we-intention to perform X or at least they ought to accept it;

3) there is a mutual belief in *g* to the effect that 1) and 2).

In clause 1) the operative agents' we-intention has as its content the bringing about of X. X might occasionally be a state that they can bring about only by themselves directly performing something else which will bring about X.³ I will not here explain (GI) further, except for later commenting briefly on we-intentions -- a central ingredient in the analysis.

One can, indeed, discern a common pattern or model for analyzing group attitudes here. This model applies to any normatively binding attitudes that a group can have. By active I mean that the group can by its own action intentionally acquire the attitude in question. These attitudes include goals (wants), intentions, and beliefs, as long as they are so acquirable. We can also speak of e.g. group fear, but it is not a normatively binding attitude in the group. I will later show how a "passive" attitude like this can be acquired and had by the group.

The general content of this model can be sketched as follows. I will use the word 'proper' to label the type of attitude in question.

1) A proper group attitude (goal, intention, belief, or even action) involves the whole group in a sense of binding the group members to the attitude (goal, intention, belief, action).

2) The relevant and required sense in which group members can thus be bound to the attitude (goal, intention, belief, action) is normative in the sense of an objective obligation to accept the attitude (goal, intention, belief, action) as applicable to all group members. The objectivity of the obligation here means its "epistemic objectivity", viz. its being treatable by actual and potential and thus future group members as something objective that does not merely consist of a believed or mutually believed obligation. In the case of operative members their acceptance of the obligation must be involved. Then we speak of their being committed to the attitude (goal, intention, belief, action). The obligation in question must be mutually believed (at least among the operative members) to exist in order for it to be -- or even begin to be -- motivationally effective.

3) Group-obligations in the sense of 2) can only come about due to the exercise of the group's relevant authority system (viz. the group's decision-making or joint intention formation system or, put slightly differently, system for jointly accepting propositions for group use).

4) An exercise of the group's authority system involves its operative members intentionally jointly agreeing to accept the attitude (goal, intention, belief, action) in question. Agreeing (thus, jointly intending) to accept conceptually entails accepting; there is no conceptual room for an intentional agreement or intention to accept and failure to accept. We may say that the operative members collectively accept, to take up the central cases of intention and belief) either "Our group's view is that p" (acceptance belief) or "Our group will do X" (intention, when 'will' is conatively rather than predictively used). Making an agreement also entails the coming about of an objective obligation to carry out what has been agreed to (and accepted). As agreement can only come about due to intentional communication of a relevant kind, it must be intentional; and there must accordingly be a mutual belief about the fact that an agreement has been effected and normally also about the content of the agreement.

The general model indeed boils down to the exercise of the group's authority system: Clauses 1), 2) and 3) entail that proper group goals centrally involve an exercise of the group's authority system. Clause 4) (together with 2)) spells out exactly what that involves. The present argument concerns the central case of a group attitude which is expressed by sentences like "Group g has goal G". This is the objective sense which amounts to more than believed or mutually believed goal, for even if some members would not have the belief that G is the group's goal, it would still have it as its goal, but yet some members are required to have the belief in question. The existence of a group

attitude in this kind of normative group-binding sense is closely related to the objectivity of the obligation involved in an agreement. If some persons make an agreement with a certain content, the agreement obligates them to satisfy the content. Other things being equal, the agreement (viz. the obligation involved in it) is valid even if the participants would temporarily forget it. The present model of group attitudes indeed involves either explicit or implicit agreement at least by the operative members, and the non-operative members must “tacitly” accept what the operative members have agreed upon (this tacit acceptance is weaker than implicit agreement, see Tuomela, 1995, Chapter 6, for clarification).

Here are three further comments on the model at hand:

a) In the terminology of Tuomela, 1995, normative group attitudes and actions in the present sense are r-notions and not s-notions. An r-notion is a rule notion, where the rule is ultimately based precisely on agreement making. An s-notion in contrast is a notion based on belief or mutual belief.

b) The positional account gives positional attitudes and actions, and such an attitude or action is precisely one in the *we-mode* as defined in the present paper (see especially note 1).

c) The present positional model of group attitudes involves collective commitments in a strong sense. Let start be commenting on collective and joint commitments in general terms (cf. Tuomela, 2000a, Chapter 2). To keep matters simple, consider a collective case with only two agents, a and b. I will concentrate on the case where they both are individuals (rather than collective agents). I will next define some of commitment notions for the case where a and b are supposed to commit themselves to perform a collective (viz. many-person) action a, such as what the operative members do according to the positional model. I start by taking social commitment (Scom) as my general logical notion concerned with a being committed to b to do something Y, taking X to denote a collective (e.g. joint) action or “project” (which could also be an institution):

- $SCom(a,b,Y)$ = agent x is socially committed to y to perform action Y;
- $Com(a,X^a)$ = a is committed to performing X^a as his part of X;
- $CoCom(\{a,b\},X)$ = a and b are collectively committed to performing X together;
- $SCom(a,b,X^a)$ = agent a is socially committed to b to do his part X^a of X as his part of it;
- $SCom(b,a,X^b)$ = agent b is socially committed to a to do his part X^b of X as his part of it;
- $SComCoCom(\{a,b\},X)$ = agents a and b are collectively and socially committed to performing X together = $CoCom(\{a,b\},X) \& SCom(a,b,X^a) \& SCom(b,a,X^b)$.

According to the last definition a and b are socially and collectively committed to performing X together if and only if both are collectively committed to performing it and are also socially

committed to each other to perform their parts of this many-person action. Note that obviously $\text{CoCom}(\{a,b\},X)$ entails the conjunction $\text{Com}(a,X^a) \ \& \ \text{Com}(b,X^b)$.

Standard joint intentions (as they will be commented on in Section III below) entail collective commitment at least in the sense of CoCom . This is simply because joint intention entails we-intention to perform or bring about the joint action in question (here a) and also, in conformity with the standard analysis of a we-intention, the intention to perform one's part. Thus, as each intention entails a corresponding commitment, we have CoCom and $\text{Com}(a,X^a) \ \& \ \text{Com}(b,X^b)$. As to SCom , its presence is not conceptually entailed from the mere assumption of sharing a plan unless agreement-making is involved (or at least I cannot prove it), although social commitment seems typically be present as a contingent, mutually believed element, when some participants share a plan for joint action. Even more typically there will be social expectations of part performances. This means that the participants expect (factually and/or normatively) the others to perform their parts and also believe that the others expect that they themselves perform their parts. Such social expectations must of course be present for a social commitment to do any motivational work, so to speak (viz. lead to relevant activities, such as sanctioning). Be normative expectations present or not, in any case the participants are entitled to such expectations on the basis of the social commitments in question.

In the case of full-blown agreement-making we have SCom , but we also have obligations backing and justifying the commitments in question (viz. in our present case CoCom , $\text{Com}(a,X^a)$, $\text{Com}(b,X^b)$, and Scom). This is the situation with normative, group-binding beliefs and intentions and other group attitudes. (More generally, even in the case where full-bodied agreement is not present, both CoCom and Scom , and hence SCom , may be based on personal obligations or norms or otherwise involve normative elements (e.g. private norms). As said, obligations in an objective, "public space" sense become involved when promises or agreements concerning the performance of the many-person action are made.)

In the case of structured groups, the operative members will be normatively committed in the sense of SCom to their tasks. However, that may take place indirectly in that they may themselves have to make relevant practical inferences to find out what specific tasks they are obligated to or are permitted to perform. What the general rules and charters (and for that matter normative social expectations) really implicate for some practical situations may not be that clear and the position holders may have to suitably "interpret" the cases at hand (viz. perform relevant particularizing inferences).

Leaving normative group-binding attitudes we now consider non-normative group attitudes. These are taken to be *shared we-attitudes*. We-attitudes are attitudes involving social beliefs. We consider a person's we-attitude, WATT, related to an attitude, ATT, which has content p. Here ATT can be a plain want, goal, intention, belief, etc. (We can, in addition, let ATT be exemplified by an action performed with a certain intention or purpose; or from a logical point of view ATT can indeed be any content.)

A we-attitude in its strongest sense, we could say *core* sense, is defined as follows for a person, say A: a) A has ATT(p) and b) he believes that also the others in the group have ATT(p) and he also c) believes (or at least is disposed to believe) that it is mutually believed (or in a weaker case plainly believed) that the members have ATT(p). Acting for a we-attitude related to ATT(p) entails acting in the right "ATT-realizing" way for the reason content p (cf. clause a)) and also acting for the social reason that the others in the group are have ATT(p) and are satisfying it or at least are disposed to satisfy it (clause b)), and also acting in part on the mutual belief in question (clause c)). For instance we can have p = "Our club house is beautifully decorated" and ATT = want. When some persons act for (because of) this we-want to have the house beautifully decorated each person has a composite full reason because of the collective end in question. This full reason "internalistically" described consists of his wanting to have the house beautifully decorated and his believing that also the others want so and -- in the core case -- his also believing that it is mutually believed that the members want to have the house beautifully decorated.

How can conditions a)-c) defining a we-attitude be justified? First, it must be said that they are obviously somewhat idealized in requiring without exception that all the members satisfy the condition. That can easily be remedied, although I will not do it here. As to a), it can be thought of as a rather obvious requirement for distributable (and distributed) group properties, which in principle must apply to all individual group members. This is a central idea when speaking of we-attitudes, since they are basically attitudes that individuals can have -- as distinct from group properties like group cohesion. It can be said that we-attitude concepts are group concepts (at least in a weak sense) which have application to both the individual and the group level. An attitude can be held either in the I-mode or in the we-mode. When it is in the we-mode it requires that the participants must hold it *qua* group members, and hence they must collectively accept it as their group attitude and must be collectively committed to its content which is "collectively available" to them (or, equivalently, is meant "for the use of the group"). In contrast, I-mode attitudes are private ones, to which the person having it is only privately committed (cf. note 2).

Concerning b), belief is an obvious element for intentionally “gluing” the attitudes together, as we are speaking of attitudes distributed among the group members. A mere aggregate of facts about people having the same attitude (say a fear) is not fully social, because it lacks an appropriate doxastic connection between the participants. This connection here basically consists in the participants’ belief that the others share the attitude; and this can be a person’s reason (or cause, as the case may be) for conforming and for his continuing to have the attitude. It accordingly seems appropriate to require b), understanding belief to be the basic cognitive notion for capturing information from the surrounding social world.

What about the mutual belief requirement c)? The basic idea is that everyone should be aware that the others are socially conscious, and are disposed to act on the social, intersubjective consciousness afforded by mutual belief. Thus, a group member should not only believe that the others have the attitude in question but should also think that the others think that he has it. Often this created experiences or believed social pressure, since the agent may also think that the others think he should (continue to) have the attitude in question. A way to defend such a loop belief derives from the assumptions that the others are relevantly similar to him - if he believes that the others have the attitude (assumption b)), he should also believe that on the basis of such similarity the others believe that he also has the attitude; and this reasoning can be replicated, resulting in mutual belief in the sense of replication or iteration (see Tuomela, 2000a, Chapter 5). There should also be similar replicable loop beliefs between all group members, at least if the group is small, and our reference point member should believe so (e.g. x believes that y believes that z believes that y has the attitude in question). This is A’s belief about a replicable loop belief, and it should be allowed his belief about such loop beliefs between other members can be wrong. Thus, one can speak of a member’s belief that there is a mutual belief that every member has the attitude in question. (At least a fixed-point characterization of mutual belief can allow that there may be beliefs about mutual beliefs and even mutual beliefs about mutual beliefs.)

In symbolic terms a we-attitude WATT that a participant has is expressed as:

$$WATT_x(p) \leftrightarrow ATT_x(p) \& B_x((y)ATT_y(p) \& MB((y)ATT_y(p))).$$

Here B stands for belief and MB for mutual belief in the group, and ATT represents some attitude. The notion of a we-attitude in general, without any qualifications, expresses sociality in the general sense and an action performed because of a we-attitude is social in a broad sense.⁴

A stronger intuitive idea related to we-attitudes is that a person has ATT in part *because* the others have ATT and this is mutually believed in the group. My above definition does not, however, require that this is invariably the case. Surely also the weaker notion can have applications even if

they are less social. The stronger notion can be expressed symbolically using $/_r$ for the reason-relation as follows:

$$WATT_x(p) \leftrightarrow ATT_x(p /_r B_x((y)ATT_y(p) \& MB((y)ATT_y(p)))).$$

Let us apply these notions to goals and beliefs. In the case of goals (G) we have, respectively, for the notion of we-goal, shared we-goal, and their reason-involving variants:

$$WG_x(p) \leftrightarrow G_x(p) \& B_x((y)G_y(p) \& MB_g((y)G_y(p)))$$

$$SWG_g(p) \leftrightarrow (x) (G_x(p) \& B_x((y)G_y(p) \& MB_g((y)G_y(p))))$$

$$WG_x(p) \leftrightarrow G_x(p /_r B_x((y)G_y(p) \& MB_g((y)G_y(p))))$$

$$SWG_g(p) \leftrightarrow (x) (G_x(p /_r B_x((y)G_y(p) \& MB_g((y)G_y(p)))))$$

In the case of belief we get respectively:

$$WB_x(p) \leftrightarrow B_x(p) \& B_x((y)B_y(p) \& MB_g(p))$$

$$SWB_g(p) \leftrightarrow (x) (B_x(p) \& B_x((y)B_y(p) \& MB_g((y)B_y(p))))),$$

and analogously for reason-involving we-belief. (Note that in the case of belief we get the reduction to

$$WB_x(p) \leftrightarrow B_x(p) \& B_x(MB_g(p))$$

if we use the standard kind of iterative analysis of mutual belief.)

The above notions of we-attitude, viz. the purely conjunctive and the reason-involving one, are somewhat strong. Weaker forms may be considered. We may thus speak of a *weak* we-attitude when the belief in clause c) only plain rather than mutual belief is used; and we speak of a doxastically *rudimentary* we-attitude when c) is not required at all. The beliefs and mutual beliefs might be taken to concern only the majority of group members and the mutual belief need not be taken to require that everyone participates in the mutual belief (see Tuomela, 2001, Chapter 4). Furthermore, when speaking of a shared we-attitude, the sharing may be complete or only partial. In the latter case one might speak of percentages and say that p per cent of the members of a collective have the we-attitude in question. When I in this paper speak of shared we-attitudes I basically speak in terms of the above standard notion (involving 100 per cent sharing), but the account can be relaxed. What seems a more reasonable requirement is that most group members share the we-attitude in question. This statistical majority criterion entails that the social practice is prevalent, supposing that the people in question indeed act on their we-attitudes.

Going in the other direction, shared we-attitudes include also strong jointness items. Thus, for instance, we-mode joint intentions can be included in them. In their case we must use collective

contents and speak of the participants' intentions to perform their parts of a joint action. This is we-mode joint intention. We can say, roughly, that a member of a collective g ("we" for him) we-intends to do X if and only if he; (i) intends to do his part of X (as his part of X), (ii) has a belief to the effect that the others also intend to perform their parts of X (and that, more generally, the other joint action opportunities for an intentional performance of X will obtain); and, furthermore, (iii) believes that there is (or will be) a mutual belief among the participating members of g to the effect that the other participants intend to perform their parts of X (and that the joint action opportunities for an intentional performance of X will obtain). Shared we-intentions of which the participants are mutually aware are called joint intentions. Also agreement-based joint intentions can be included here simply by assuming that an agreement to perform X jointly has been made and that the participants hence are jointly obligated to performing X . (There are several other "variations of the theme we-intention" for which I refer the reader to Tuomela, 2000a, Chapter 2, and 2000c.)

It is clear that normative, group-binding beliefs and beliefs in the shared we-belief sense are quite different. As to our example cases, it can now be said that cases g) and i), and often also h), are best analyzed in terms of the we-belief account (although h) might also and in addition qualify as a proper group belief). Indeed, in many cases a proper group belief that p and a shared we-belief that p will both exist.

Here is a small example, which also shows the difference between our two kinds of group belief in a kind of conflict case. Consider a dyad consisting of a and b and assume both for the case of normative and nonnormative group belief:

- 1) $B_a(p)$
- 2) $\neg B_b(p)$ (or possibly even $B_b(\neg p)$)

These beliefs 1) and 2) are I-mode beliefs or, possibly, acceptances. Now, we may in addition, have in the normative case (but not in the nonnormative case, except when we are dealing with the peculiar and problematic case where 1) and 2) really represent beliefs):

- 3) a and b jointly accept p

Note that in the normative case 3) must be a we-mode joint acceptance belief (and in the mentioned peculiar nonnormative case it might also be one).

We then have for the normative case, assuming that the group's authority system is at work,

- 4) $B_g(p)$,

which says that group has the normative, group-binding belief that p .

But in the the nonnormative case construed here as an aggregative, shared we-belief account, which involves only I-mode beliefs, we get

5) $\neg B_g(p)$

Nevertheless, it is possible (at least when 3) happens to hold for both the normative and the normative case) that within both approaches

6) B_g (both a and b jointly accept p as the belief of g).

However, one still partially reconcile the two approaches and form a disjunctive combined account. I will not discuss this much further here but only give a partial summary of the consequences of the accepted positional account of proper group beliefs and of the we-belief account in the case of other group beliefs. Thus, the following features of the final, *combined account of group beliefs* can be stressed (for detailed discussion see Tuomela, 1995, Chapter 7, on which source I here draw):

1) Groups (e.g., formal and organized groups) in general can have normative group beliefs such as the belief that something p in the wide sense of accepting p.

2) A group's having a normative group belief that p is compatible in the extreme (or spurious) case with even without all or even any of the members' personally (viz. in the I-mode) believing that p (and in principle also with their believing that not-p). However, it is relevant to the genuineness and firmness of the group's belief that p to what extent its members personally believe and, indeed, we-believe that p.

3) Any member of a group has the *prima facie* obligation to accept (in the discussed wide sense) p *qua* a member of the group (viz. in the we-mode), if the group has the normative group belief that p; but such a member is under no obligation to accept or believe that p in the I-mode.

4) Gallup-poll type of statistical investigations of the relative frequencies of people sharing a we-belief are at least theoretically possible. However, the empirical investigation of normative group beliefs will have to be quite different, the operative and non-operative members here obviously playing a quite different role.

5) A group can have a normative group belief only if it incorporates either an explicit or an implicit group-authority system on which the collective or joint acceptance of views relies.

6) Groups can have non-normative beliefs in the sense of shared we-beliefs even when they have no normative beliefs. This is the case when no group-authority system is at work to create group-commitments involving the group as a whole. A normative group belief is the more genuine the more widely the group members' corresponding personal (viz. I-mode) we-belief is shared.

Mutatis mutandis, the points 1)-6) apply to group goals as well.

Let me still emphasize that in terms of our we-mode/I-mode distinction, normative group-binding group beliefs are based on operative members' we-mode beliefs (or, as I also say, positional beliefs), whereas we-beliefs in the non-normative sense typically are I-mode beliefs (but also they can be we-mode beliefs).

IV COOPERATION

I will briefly consider cooperation here, as the central notion is a we-mode notion, even if normally in e.g. game theory and DAI cooperation is treated as an I-mode notion. I will now summarize my central arguments for the we-mode character of cooperation in the standard sense. The arguments show that mutual many-person cooperation must be based on the shared we-mode collective goal that the participants have accepted (see Tuomela, 2000a, Chapter 4):

1) a) The conceptual argument is simply that the very notion of full-blown cooperation is taken to rely on a collective goal in about the sense of my analysis of the notion of a we-mode collective goal. b) Linguistic evidence also supports this view, even if one should be cautious about dictionaries for purposes of conceptual analysis. Thus, according to Collins' dictionary, to cooperate is to act together for a purpose, and, according to Webster's dictionary, to cooperate is to work with another or others to a common end. Note, however, that an expression such as 'common end' still can be taken to cover also shared I-mode goals, resulting in I-mode cooperation. However, this is no problem for my argument as long as genuine cases of we-mode cooperation exist.

2) One may view cooperation from a group's point of view: the group has achieving something as its goal and intentionally acts to achieve it; and, conversely, when the group acts intentionally there must be a goal of some kind involved. The central point here is that a we-mode collective goal offers more stability – in a psychological, epistemic, and normative sense – than do I-mode goals, even share ones. The argument from group analogy entails this: If a) the group members have the same goal G, b) collectively accept to achieve this goal for their collective use, and c) are collectively committed to achieving G and act so as to realize this commitment, then the goal G is a we-mode goal, viz., a collective goal satisfying the Collectivity Condition (CC). The conditions a)–c) can be regarded as satisfied in the case of full-blown cooperation involving collective or joint activity. The members' collective commitment to G is central here in showing that a shared we-mode goal and not only an I-mode goal must be present. Consider an example: I am mowing the lawn and you are planting flowers in a garden. This could be I-mode cooperation with different private goals (to have the lawn mowed and the flowers planted), it could be I-mode

cooperation towards a shared I-mode goal (cleaning up the garden), or it could be a case of we-mode cooperation. In the first case we separately perform our activities and are free to change our minds about our tasks without the other's criticism. In the last case we collectively accept to clean up the garden and are accordingly collectively committed to doing so. Our collective action of cleaning up the garden is a cooperative joint action. The upshot of the present argument is that the adoption of the group perspective warrants the claim that there must be cases of g-cooperation. The categorical premise for the necessity argument is that a group perspective not only is often adopted but must be adopted by such social group beings as human beings basically are.

3) Related closely to the second argument, we have the following normative argument: Actual life abounds with cases of cooperation in which the participants take themselves to be collectively committed to cooperative action and accordingly tend to think partly in normative ways such as "I will participate because I ought to do my part of our joint project". This collective commitment, which is stronger than aggregated private commitment, indicates the presence of a collective goal and we-mode cooperation, for the norm is one related to the group context in question.

4) Instrumental argument: Shared collective goals tend to work or function better, e.g., for achieving coordination and stability than shared private goals (and shared compatible private goals). In some cases – such as in the case of games of "pure coordination" and the "centipede" – the required kind of coordination cannot optimally be achieved without the participants sharing a collective goal. (Various aspects of this argument are considered in Tuomela, Chapters 2–4, 6, 9, 11, and 12.)

For lack of space, I cannot here discuss cooperation in detail. To make a long story short, I will give a simple semi-logical analysis. It uses the operator DO for collective action and requires mutual belief also about the collective action:

$$\text{MCOOP}(x,y,p) \leftrightarrow \text{CGoal}(x,y,p) \ \& \ \text{DO}(x,y,p) \ \& \ \text{MB}(\text{CGoal}(x,y,p) \ \& \ \text{DO}(x,y,p)).$$

CGoal(x,y,p) means that x and y have a group-mode content p as their collective goal in a we-mode sense. The we-mode sense will satisfy the so-called Collectivity Condition, according to which out of "conceptual" necessity, or rather necessity based on the participants collectively accepting the goal as their collective goal, the goal (here p) will be satisfied for all participants if it is satisfied for any for them (see Tuomela, 2000a, Chapter 2, and especially Miller and Tuomela, 2001, for further requirements). Here p can be taken to have the "canonical" form "We perform X together" or "We bring about a state X together" (with the appropriate verb tenses depending on the context). Understood in this way, p will express a divided goal, and furthermore, it will be taken below to satisfy the Collectivity Condition. As before, we will accept the requirement that each agent must contribute to the collective action. In symbols:

$DO(x,y,p) \rightarrow DO(x,p_x) \& DO(y,p_y)$.

Here p_x and p_y are the respective shares (or part-actions or contributions of the agents x and y . Collective action need not be successful, but does involve acting purposefully to realize the goal of the action.

In my simplified formalism the notion of a collective goal comes out as follows in view of the earlier analysis of reason-based we-attitudes (here we-intentions I_{we}^r) of Section III:

$CGoal(x,y,p)$ if and only if $(I_{we,x}^r(p) \& I_{we,y}^r(p))$ and accordingly if and only if $(I_x(p /_r I_E(p) \& MB(I_E(p))) \& I_y(p /_r I_E(p) \& MB(I_E(p)))) \& CC(p)$,

where $I_E(p)$ means everyone's intending that p , $CC(p)$ the satisfaction of the strong collectivity condition by p , and MB means mutual belief.

Mutual cooperation with a *cooperative attitude* ($MCOOP_f$) can here be symbolized simply by writing WDO to mean willingly (viz. not reluctantly) performed collective action:

$MCOOP_f(x,y,p)$ if and only if $CGoal(x,y,p) \& WDO(x,y,p) \& MB(CGoal(x,y,p) \& WDO(x,y,p))$. I want to emphasize that in my formal description the collective goal can be merely the jointness-element or action-goal involved in the cooperative activity itself (rather than being a separate end-goal).

V OTHER SOCIAL MACRO CONCEPTS

Now that we have discussed both normative, group-binding and non-normative group attitudes and actions it can be asked what other social macro properties are there to be accounted for. As before I will restrict my treatment to the man-made aspects of the social world, viz. to those aspects of the social world that the participants in social life themselves in some sense construct or can construct. This means that I will not discuss such things as social features which come about as a matter of causal fact without nobody voluntarily bringing about them. I suppose that social psychological relationships like love and we-hate or we-fear may well belong to such features of the social world. There also seem to be social relations which are at least in part man-made in a "constructivist" sense. Thus kinship relations seem in part objective and in part socially constructed (cf. the concept of gender and the norms related to patriarchy or matriarchy). Social relationships such as those of production, as Karl Marx discussed them, might seem not to be socially constructed in the sense meant here. However, I doubt that. For instance, the capitalist and the worker are in part in a norm-governed relationship, and norms are man made. (However, a proper discussion of the economic relations of production, etc. must be left for another occasion.)

A social theorist's concepts such as group cohesion, suicide rate, stratum, or class are not collective man-made notions in the sense meant here, and thus they fall outside the present investigation.

What is of more interest, there are unintended consequences of intentional action which amount to social macro properties. Thus, the fact that under suitable market conditions supply and demand of some goods will optimally meet one another is one such unintended macro thing, if we are to believe Walras and other classics of economics. This kind of phenomena do not depend on collective intentionality but are aggregative outcomes of I-mode intentional action. However, as thesis 9) argues, also unintended and unforeseen consequences can have a role in the dynamic picture, because they will affect the collective attitudes underlying social practices and also the underlying institutional and other conditions ("collective action opportunities") that recurrent collective activities rely on. More precisely, this kind of unintended results (as a kind of goals), via feeding back information from those results, may repetitively affect the agents collective attitudes and hence their relevant activities. The feedback thus can occur without the agents being aware of it, and thus we can have functional, self-reinforcing invisible hand processes in the strong sense Adam Smith seems to have conceived of them: "[the individual] generally, indeed, neither intends to promote the publick interest, nor knows how much he is promoting it ... he intends only his own gain ... and he is led by an invisible hand to promote an end which was not part of his intention" (Smith, 1982/1776, IV, ii, 9). The account of the dynamics of social practices developed in Tuomela (1995) and enriched and improved in Balzer and Tuomela (2001) allows for unintended ends and results of the above strong Smithian kind.

The above remarks purported to show that there seem to be some social phenomena which fall outside the scope of the Collective Acceptance account and thus, so to speak, beyond the governance of collective intentionality. However, the issues are not that clear, and further research is needed show which social phenomena totally fall outside the constructivist scope. On the positive side, social practices such as customs and traditions clearly are man-made and man-maintained, and so are social institutions. The Collective Acceptance model is able to account for all of these (see Tuomela, 2001). Briefly, social practices (repeated collective social actions in my terminology) are in their core sense argued to depend on collective intentionality in the sense that they are performed because of shared we-attitudes so that either a we-intentions or at least we-beliefs are involved. As collective acceptance amounts to the group members' coming to hold and holding a relevant we-attitude, specifically a we-intention or a we-belief, we can see how the Collective Acceptance account accounts for social practices. Consider next social institutions in the wide, general sense

covering e.g. money, language, property right, and so on. Here is a summary of my basic analysis of such “standard” social institutions:

(SI) A generic sentence *s* expresses a *social institution* (in the “standard” sense) in a collective *g* if and only if

1) *s* expresses or entails the existence of a social practice (or a system of interconnected social practices) and a norm (or a system of interconnected norms) for *g*, such that the social practice is performed at least in part because of the (content of the) norm;

2) the members of *g* collectively accept *s* for *g* with collective commitment; here it is assumed that collective acceptance for the group entails and is entailed by the correct assertability of *s*.

In this analysis the notion of norm is a social macro property, but according to the Collective Acceptance model it is performatively generated by either by the members of group *g* or by members of some other groups that may have existed far in the past. Furthermore, norm-obeying is action, typically we-mode action, which does not involve spooky ontology.

It can be remarked here that in the case of structured collectives the operative members will typically create the rule-based institutions. Such institutions often serve to solve collective action dilemmas. In such cases there is a conflict between the collective’s (or the “public”) interest and the preferences of the individual members of the collective, viz. there is a conflict between we-mode and I-mode preferences.

How are properties collectives related to the properties and relationships between their members? What is the connection between man-made macro properties of the social world and the properties and interrelations between human beings in living the social world? What in general is relationships between the macro level and the meso level in the social world? (I define my concepts so that the meso level includes the micro level of individuals and their properties as its part.) I submit that the general answer to all these three, increasingly more general problems is that the sought after ontic relationship is supervenience, indeed strong supervenience. Let me now briefly consider these issues.

Supervenience is a relation of determination that purportedly holds between two ontic levels, the base level and the supervenient level. Thus, in the social case the base level is the meso cum micro level. Base level phenomena are supposed to determine the phenomena at the supervenient level, and corresponding to any changes at the supervenient level there are supposed to be changes at the base level which account for the former changes. Accordingly, it can be said that basically supervenience has the form of a (mathematical) function: the facts on the supervenient level are in a functional relation to facts on the base level. This is the received story, so to speak. But one can, at

least to some extent, discuss the issues involved here without taking explicit stand on ontological issues and discuss the issues in a linguistic fashion. After such a linguistic turn one can of course come back to ontological issues. Below I will mainly operate linguistically, to keep issues clearer, although my main thesis is, after all, ontological.

Recall that we distinguished two different problems in accounting for the conceptual nature of attitudes and actions by collectives. First, the notion of the right normative and social circumstances C must be an understood notion; secondly, a structural account of the relationships between the holistic, collective level and the meso level or, to use another term, jointness level must also be given. Here we will discuss the second problem but from the new angle of supervenience, so to speak, and rely on the Collective Acceptance model in the form requiring that the group is exercising its authority system (but will not explicitly mention it). The fact that that the macro properties, group attitudes and actions under discussion here, are ultimately based on collective acceptance means that the supervenient level is “conventionally” connected to the based level, to use a standard term. The conventionality here is simply that in the last analysis the connection is taken to be man-made.

It will be argued below that group attitudes and actions supervene on the attitudes and actions of their operative members (or non-member representatives), and especially on the joint attitudes and actions, respectively. We will accordingly arrive at the result that group actions and attitudes are constituted by their members’ relevant actions and attitudes.

The literature on supervenience has concentrated on the supervenience of the mental on the material roughly in the sense of the following two claims: First, mental features are physically "embodied" or "grounded" by means of some physical counterpart features, and, secondly, if two persons agree in all their physical features that determines that they also agree in their mental features (cf. Kim, 1987, for discussion).

Obviously, both social rules and proper social norms can give structure to social groups (see Tuomela, 1995, Chapter 4, for a detailed account). For our present purposes we can simply lump together all these norms and call them the norms R. So we will be dealing with a group, say g, which is characterized by a possibly long conjunction R of relevant norms.

The *thesis of supervenience* which will be discussed here is that the attitudes and actions by the collective will supervene on the actions and joint actions performed by the (operative) members of g. Supervenience will involve two claims: *the "embodiment" claim* and *the determination claim*. The embodiment claim says that whenever g has an attitude or does something, let us use the predicate F for this, then, respectively, the attitudes or actions will have to be involved: Collectives without acting members (position-holders, and what have you) cannot have attitudes and cannot act.

The most central argument for this claim goes by means of examples. For instance a business firm can have a belief and can negotiate a deal only via some of its members or representatives beliefs and actions. As to the determination claim, it says that, respectively, suitable attitudes or actions (or consequences of them) by the members of the collective will "conceptually" (semantically, conventionally, legally, or in some other relevant non-causal sense) determine the collective's attitude or action (viz., what the collective does). The general argument for this claim can be given by help of a replica of g . Let us concentrate on the action case and construe a replica of g , say g' , which is governed by the same norms, viz., its norm conjunction R' coincides with R , and in which, furthermore the members - assumed to follow the norms and not to break them - act similarly in those cases where the norms don't dictate what to do. Then it seems compelling to think, given that g and g' also share their legal, social, physical and other relevant circumstances, that these two collectives must act similarly. We may also give this argument a slightly different twist and speak of the justification and explanation of the task behavior (and role behavior) of the members of g and g' . Suppose we partially explain director A 's signing a document in some appropriate circumstances C by saying that firm g is buying a piece of land via his action. Then, other things being equal, we must also similarly explain the signing by counterpart-director A' of the firm g' by saying that g is thus buying land. (It would of course be crazy to think that g' is selling land or donating to charity or doing something else.)

It will assumed in our technical analysis that the fragment of natural language we need for describing collective action can be partially regimented in a suitable formal language and that resources for intensional operators, such as necessity, are available (recall the large body of literature related to Davidson's program for formalizing action sentences). For simplicity's sake I will concentrate on the case of action and draw on my treatment in Tuomela, 1995, Chapter 5. But it is easy to see that the formal treatment works quite similarly for group attitudes as it does for group actions. To proceed with the action case, the basic, primitive action predicates are assumed to be available *in principle* (although not necessarily in any existing human language) and the same goes also for the compound predicates that can be formed out of them in the language. Let now A stand for the set of (primitive) macro-level predicates applicable to collectives - much as folk sociology (and social psychology) has it. I will below speak of A as it were a set of action predicates, but it can as well consist of goals, intentions, beliefs and whatever attitudes one can attribute to collectives. Next, let B stand for the (both primitive and compound) base level predicates and in the action case it will stand for action predicates applicable to single action and joint action by individuals, especially members of collectives. No doubt there is much overlap between these two sets of predicates. We take each of them to be generated or generatable from a finite basis of primitive predicates (cf. Davidson's learnability of language argument as backing this). We will

assume that there are maximal finite conjunctions of all the primitive non-compound atomic predicates, unnegated or negated. We can take these primitive non-compound predicates to represent non-generated, *basic* single-agent or multi-agent actions. We will assume that the language of analysis will contain the resources for second-order quantification as well as a suitable necessity operator, say N , of a standard kind expressing, generally speaking, non-factual, such as legal, conventional, semantic, etc., necessity. One way to define this notion of necessity is to do it metalinguistically and by reference to the norms R defining and regulating the collective in question (cf. Tuomela, 1977, Chapter 9).

Using more or less a Davidsonian formalization of action sentences with event ontology, I propose the following, somewhat simplified definition schema for actional supervenience of A on B , relativized to the appropriate social-normative circumstances C (cf. the somewhat different notions in Kim, 1984, p. 165, Tuomela, 1984, p. 24, and 1985, p. 111):

(AS) A supervenes on B , given C if and only if in C , $N(x) (F \text{ in } A) (F(x) \rightarrow (E y_1) \dots (E y_m) (E G_1 \text{ in } B) \dots (E G_m \text{ in } B) (G_1(y_1) \& \dots \& G_m(y_m) \& N(z_1) \dots (z_m) (G_1(z_1) \& \dots \& G_m(z_m) \rightarrow (E v) F(v))))$, for some m .

Here A may be called the *supervenient* family and B the *base* family of predicates. In this definition the small letter variables x, y, z, v range over token events or states, In the case of action that I will concentrate on here they thus range over action tokens, and ' $F(x)$ ' reads 'action token x is an F 'ing performed by the collective' (explicit mention of the group has been suppressed above), ' $G_1(y_1)$ ' reads ' y_1 is an action token of G_1 'ing', and so on. The variable F and the variables G_i are assumed to be replaceable by the suitably regimented counterparts (predicate constants) of natural language action predicates, such as 'declares war', 'signs a declaration', 'ratifies', and so on. The subscript m is a variable governed by an existential quantifier whose scope is the entire analysans. Thus, assuming, as is standard, that one agent performs only one action token in this context, our analysans says that, for each collective predicate F in A (such as 'declares war'), and for each token of that predicate, there are m agents in the group performing G_1, \dots, G_m , for some m , and that these performances jointly generate or bring about x , viz., the performance of F by the group. Note that, corresponding to different tokens of F , different G_i s may be involved. Thus declaration of war may take place in several different ways, viz., be determined by tokens of different kinds of joint action.

As before, we need not be able to say that these agents jointly do F but only that they jointly do something which will bring about F (and this something could be a joint action G , nonidentical with F). However, we may decide to speak of joint actions whenever the agents jointly bring about some result, and so we would say also in this case that the agents jointly performed F (and not only G). The generation in question here takes place mainly in virtue of the norms R (which in our

example will or may involve, e.g., the constitutional legal rules for a state's declaration of war) in those circumstances, and while it may contain factual aspects it is primarily conventional, or “quasi-conceptual” and man-made, in the sense of the Collective Acceptance account. (Note that the above definition uses only monadic action predicates, and is in this sense simplified.)

In the case of group attitudes the predicate F will of course represent an attitude and the variables to which it applies will be states of g . In the based we will correspondingly have joint attitude predicates.⁵

Given our above notion of supervenience and given a couple of additional assumptions to be mentioned below, the following result follows (cf. the analogous result based on a very different kind of proof in Kim's, 1984, system):

(#) If A supervenes on B , given C , then, in C for each predicate F in A there are maximal predicates H_1, \dots, H_k in the base B such that $N(x)(F(x) \leftrightarrow (E y_1) \dots (E y_k) (H_1(y_1) \& \dots \& H_k(y_k)))$, viz., a collective's actions have a necessary coextension in the base level domain.

For the proof of (#) see Tuomela (1995), Chapter 5. We can now claim that (AS) and (#) jointly show the truth of the following claim (for a collective G with m operative agents):

(##) If A supervenes on B , given C , then for every predicate F in A it is true in the circumstances C that the collective satisfies F if and only if some operative members (or representatives, more generally) bring about the satisfaction of F .

The truth of the sufficiency-part of (##) follows from (AS) and the discussion following it. It should, however, be noted that (AS) may (and probably often will) become satisfied by means of only a minimal jointness level basis for the group action. In such cases the m jointness level token events or states satisfying it are necessary for it; but we should allow that there be more than m operative members (jointly) performing something which makes F satisfied. ((##) strictly speaking requires that only such "minimal" sets of operative members, satisfying (AS), be involved. See Tuomela, 1995, Chapter 5 for a discussion and proof of (##).)

One especially interesting relation in the context of group action and attitudes is the one which allows us to say that a collective's action and attitudes are constituted by its (operative) members' actions and attitudes, respectively. Consider this:

(CON) A group's satisfying F (viz., an action or attribute predicate) in the social-normative circumstances C was *constituted* by some members of the group satisfying relevant predicates in the base set B if and only if

- a) A supervenes on B , but B does not supervene on A , given C ;

b) the analysis of (AS) was satisfied by some even or state tokens, i.e. in addition to the token, say x , of F there were tokens y_1, \dots, y_m of the respective kinds G_1, \dots, G_m , in C ; and the main clause of the analysis, viz., the one following the qualification 'in C ', would not have been so satisfied had C not obtained.

Clause a) states the basic connection to supervenience. Note that this analysis does not specify the basis of the constitution more than existentially (cf. (AS)), and we shall refer to this feature below. Another general matter to be emphasized here is the issue of asymmetry. The notion of supervenience is not asymmetric if merely the formal aspects of (AS) are taken into consideration (nor is it symmetric). But if we interpret the predicates in A as ones to be applied to groups in this context, then probably nobody would have reason to defend the claim that the jointness level and individualistic predicates in B are grounded in group properties. For, as is obvious, many different combinations of individual token events or states may, so to speak, realize a collective property, but surely not *vice versa*. And this argument may be taken as sufficient to establish asymmetry, which is surely needed. But I have anyhow chosen to explicitly put in asymmetry.

Clause b) comes to say, in conjunction with a), that supervenience is instantiated in the particular case at hand, and it does it by saying that the group property action had a meso level basis. This clause also requires the rather obvious thing that supervenience holds in part in virtue of the social-normative circumstances C , assumed to satisfy all the constitutive and regulative norms characterizing the group in question. And, in any case, it must be kept in mind that the connection between the macro (group) level and the meso level is conventional in nature. What our analysis shows is that we need not after all postulate macro-level entities and properties in an ontic sense. Note that my results above are compatible with postulating holistic ontology as long as supervenience holds, but here I am claiming that it is not needed.

To be sure, I have not discussed more than group attitudes and actions above and claimed the above kind of lower level ontic constitution for them. But as far as I can see only the case of institutions (or other macro-level items) can provide an obstacle to my supervenience program. However, the case of institutions, as analyzed by (SI), was already commented on. We saw that apart for positional, we-mode action and norm-referring activities not much more is involved from a conceptual point of view, and the just-mentioned notions are not spooky notions. Furthermore, they are notions attributed to individuals rather than groups, even if an institution by itself is a group level phenomenon. is to be seen as an ontic reconstruction program, which need not be seen as reductive.⁶ Furthermore, it relies on holistic concepts – our ordinary common sense concepts in a somewhat sharpened form.

It is worth emphasizing that my program which I have tried to argue for in detail in the case of group properties and in less detail for other macro-level entities properties, or alleged entities and properties, is compatible with naturalism of a kind. For it is possible to argue that mental states and social psychological relations occurring on the meso and micro levels in my system themselves supervene on non-psychological and non-intentional notions, and that would give a rather strong kind of naturalism. Even if intentionality would be needed in the last analysis as an element of the world's ultimate ontic furniture, that would still be naturalistic enough. Suppose best explaining theories of science ideally determine the world's ontic elements. Then, if intentionality in some ontic guise will be present, so be it (and so it had better be).

Notes

1) Thus, not only theorizing in the various social sciences are included but so are also such fields of Artificial Intelligence (AI) as Distributed AI (DAI) and the theory of Multi-Agent Systems (MAS). DAI studies the social side of computer systems and includes various well-known areas ranging from human-computer interaction, computer-supported cooperative work, organizational processing, and distributed problem solving to the simulation of social systems (cf. Tuomela, 1996).

2) I have given analyses of the notions of we-mode and I-mode in different works and earlier spoken also of the positional mode or the group-mode instead of the we-mode (see Tuomela, 2000, 2001, and Miller and Tuomela, 2001, for recent characterizations). I will here generalize the account of we-mode, I-mode, and private mode goals in Miller and Tuomela (2001) to cover any kind of attitude. The private mode is a special case of the I-mode, viz. a case which denies dependence on any group and not only the reference point group g .

Here are my definitions of the we-mode and the I-mode in their "standard" senses:

a) Agent A , a member of group g , has a certain attitude ATT with content p in the *we-mode* relative to a group g in a certain situation C if and only if ATT with content p (thus the sentence s expressing it) has been collectively accepted in g as g 's attitude, and A is functioning (viz. experiencing, thinking, and/or acting) *qua* a group member of g and is collectively ATT -committed to content p at least in part for g (viz. for the benefit and use of g) in C . (Here g can be either an unstructured or a structured group.)

b) Agent A , a member of group g , has a certain attitude ATT with content p in the (plain) *I-mode* in a certain situation C if and only if A is not functioning *qua* a member of g and is privately

ATT-committed to content p at least in part for himself in C. (Here g can be either an unstructured or a structured group. What is analyzed here is simply acting as a private person relative to group g but possibly not relative to some other group.)

c) Agent A has a certain attitude ATT with content p in the *private mode* in a certain situation C if and only if A is not functioning *qua* a member of any group, and is privately ATT-committed to content p only for himself in C.

Acting, or more generally functioning, *qua* a group member is a notion with several aspects, of which only one is acting in one's group position. Let me therefore here summarize the account of this notion I have given in Tuomela (2002).

In any group it may be possible to perform freely chosen actions *qua* a group member provided these actions are within the realm of concern of the group, viz. belongs to topics which are of concern or are of significance for the group as socially accepted by the group (as a normative, group-binding group acceptance or through the we-acceptances by the group members or their majority). The non-normative acceptance or belief here could be of the form of a we-acceptance within the group, g: Ideally, everyone accepts T to be a topic of concern for g, and believes that everyone so accepts and also believes that this is mutually believed in g. So we get a notion of group concern:

Topic T is within the realm of group g's concern if and only if T is we-accepted to be in the group g's realm of concern. Group g's realm of concern C consists of a set of topics $\{T_1, \dots, T_m\}$. A topic T_i which is within the realm of g's concern consists of a set of actions (types) X_j ; let us call their set X_i .

Now considering the union of the set of actions X_i , viz. $\bigcup_i X_i$, we can classify those actions as follows from the point of view of the idea of acting *qua* a group member.

The general case is that of a structured groups. In their case we can classify actions falling within the realm of actions within $\bigcup_i X_i$ as follows:

1) positional actions (related to a group position), which include i) actions (tasks) that the position holder in question *ought to* perform, perhaps in a special way, in certain circumstances and ii) actions that he *may* (is permitted to) perform in some circumstances;

2) actions which the general (viz. not position-specific) group norms (r-norms or s-norms) as well as moral norms, laws (the like regulative rules) as well as group standards require or allow;

3) actions and joint actions that do not, or at least need not, belong to classes 1) or 2) and which are based on situational agreement making and other agreement making which has not been

codified in the task-right system of *g* or the group norms of *g* but which still are consistent with actions in 1) and 2);

4) freely chosen actions or activities (and possibly joint actions), which include actions and activities not within classes 1)-3) which, although not incompatible with them, still are actions within the realm of concern of *g* and collectively accepted (or acceptable) by the members of *g* as such actions.

Functioning as a group member in the positional case, viz. in a structured group, is equivalent to acting in one of the senses 1) – 4). Obviously, one can be a group member without *always* acting as a group member. Actions in 1) are of course typical positional actions that accordingly qualify as acting *qua* a member of *g* in one's position. Subclass ii) of 1) thus consists of actions that the holder of position may choose from. (The task-right system specifying i) and ii) may contain r-norms and/or s-norms.) However, classes 2)-4) also can occur in the positional case and also in other cases. In the case of unstructured groups, class 1) is empty.

The following theses about the relationships between the we-mode and the I-mode are defensible (cf. Tuomela, 2000, Chapter 2, for a refined, strong notion of we-mode):

T1) I-mode does not entail we-mode.

T2) We-mode does not entail I-mode.

T3) An action or an attitude can be simultaneously performed (action) or had (attitude) in the I-mode and in the we-mode.

T4) Agreement making entails we-mode intention to fulfil the agreement

3) To be more precise, we may accordingly require this:

##) A_1, \dots, A_m jointly accept the we-intention to bring about *X* in *C* for *G* if and only if in *C* there is an action *Y* (possibly identical with *X*) such that these agents have the group-intention to perform *Y* with the intention of thereby generating the group *G*'s bringing about *X* in *C*.

If a social group *g*, having exercised its authority system in intention formation, has the group intention to perform *X* then each member of *g* will be involved: The operative members must we-intend (or actually to we-intend or be disposed to we-intend) to perform their parts of *X*, or, if they do not have such parts, to contribute to *X*. This is a "positional" or we-mode we-intention rather than their personal I-mode intention. In turn, the non-operative members - even those who do not endorse the group intention - ought to tacitly or weakly accept the fact that the operative members have formed the group intention to perform *X* for *G*.

4) In Balzer and Tuomela (1997) a logical theory of shared we-attitudes is developed. In it, we use two different characterizations or explicates of we-attitudes which initially can be formulated as WATT in the text – but with the addition of the mention “joint action opportunities”. Of these explicates the first one is called the “direct” characterization, and it employs the standard “iterative” characterization of mutual belief. The second one is a “fixed point” characterization, in which the essential point is to give a the mutual belief a fixed point treatment. While I refer the reader to the mentioned work for a proper discussion (and note the formulations in that paper are primarily geared to the I-mode case), let me here just indicated in terms of an example related to we-mode joint intention how the characterizations essentially come out. We are dealing with a case with the two agents *i* and *j* who jointly intend, viz. share the we-mode we-intention, to play the Kreutzer sonata together.

The bare bones of the *direct* characterization (viz. *cad()*) are simply as follows:

$\text{cad}(\langle i, j \rangle, \text{joint intention to play the Kreutzer sonata together}) \leftrightarrow$

i and *j* have the we-intention to play the Kreutzer sonata together \leftrightarrow

a) *i* has the intention to participate in playing the Kreutzer sonata together with *j*, and

b) *j* has the intention to participate in playing the Kreutzer sonata together with *i*, and

c) *i* believes that the joint action opportunities for their playing the Kreutzer sonata together obtain and that there is mutual belief among *i* and *j* that both have the intention to participate, and the joint action opportunities are satisfied; and

d) *j* believes that the joint action opportunities for their playing the Kreutzer sonata together obtain and that there is mutual belief among *i* and *j* that both have the intention to participate, and that the joint action opportunities are satisfied.

The *fixed point* characterization (*caf()*) is the following:

$\text{caf}(\langle i, j \rangle, \text{joint intention to play the Kreutzer sonata together}) \leftrightarrow$

i and *j* both have the we-intention to play the Kreutzer sonata together \leftrightarrow

a) *i* has the intention to participate in playing the Kreutzer sonata together with *j*, and

b) *j* has the intention to participate in playing the Kreutzer sonata together with *i*, and

c) *i* believes that the joint action opportunities for their playing the Kreutzer sonata together obtain and that *i* and *j* have the we-intention to play the Kreutzer sonata together with the other one and that *j* also believes analogously; and

d) j believes that the joint action opportunities for their playing the Kreutzer sonata together obtain and that i and j have the we-intention to play the Kreutzer sonata together with the other one and that i also believes analogously.

It is proved in Balzer and Tuomela (1997), Theorem 7, p. that the direct characterization and the fixed point characterization are equivalent.

5) My notion of supervenience resembles Kim's (1984) notion of strong supervenience in many central respects. One major point of difference, however, is that my approach is linguistic, while Kim' operates with nonlinguistic properties in a Boolean algebra (and with problematic property formation principles like the complement-operation). Another difference is that in (AS) the supervening predicates apply to entities different from those to which the base predicates apply

6) The ontic content of groups can be regarded as consisting of past, present, and future individuals who are position holders in the group. So we can say that these individuals make up the group. Thus the ontology would be an open sequence $\langle i_1, i_2, \dots, i_n, i_{n+1}, \dots \rangle$, where n represents the present time ("now"). Conceptually my view is holistic.

References

Balzer, W. and Tuomela, R., 1997, 'A Fixed Point Approach to Collective Attitudes'. In: Holmström-Hintikka, G. and Tuomela, R. (eds.), *Contemporary Action Theory II*, Kluwer Academic Publishers, Dordrecht and Boston, pp. 115-142

Balzer, W. and Tuomela, R., 2001, 'Collective Intentions and the Maintenance of Social Practices', forthcoming in *Autonomous Agents and Multi-Agent Systems*

Cohen, J., 1992, *An Essay on Belief and Acceptance*, Oxford University Press, Oxford

Kim, J., 1984, 'Concepts of Supervenience', *Philosophy and Phenomenological Research XIV*, 153-174

Kim, J., 1987, ' "Strong" and "Global" Supervenience Revisited,' *Philosophy and Phenomenological Research 48*, 315-326

Miller, K. and Tuomela, R., 2001, 'What are Collective Goals?', ms

Smith, A., (1982/1776), *An Inquiry into the Nature and Causes of the Wealth of Nations*, Penguin Classics, Harmondsworth

Tuomela, R., 1977, *Human Action and Its Explanation*, Reidel, Dordrecht and Boston

- Tuomela, R., 1984, *A Theory of Social Action*, Synthese Library, Reidel Publishing Company, Dordrecht and Boston
- Tuomela, R., 1985, *Science, Action, and Reality*. D. Reidel Publishing Company, Dordrecht, Boston, and Lancaster
- Tuomela, R., 1992, 'Group Beliefs', *Synthese* 91, 285-318
- Tuomela, R., 1995, *The Importance of Us: A Philosophical Study of Basic Social Notions*, Stanford University Press
- Tuomela, R., 1996, 'Philosophy and Distributed Artificial Intelligence: The Case of Joint Intention and Koint Action', in Jennings, N. and O'Hare, G. (eds.), *Foundations of Distributed Artificial Intelligence*, Wiley, New York, pp. 487-503
- Tuomela, R., 2000a, *Cooperation: A Philosophical Study*, Philosophical Studies Series, Kluwer Academic Publishers, Dordrecht and Boston
- Tuomela, R., 2000b, 'Belief versus Acceptance', *Philosophical Explorations* 2, 122-137
- Tuomela, R., 2000c, 'Collective and Joint Intention', *Mind and Society* 1, 39-69
- Tuomela, R., 2001, *Collective Intentionality: Social Practices and Institutions*, book ms
- Tuomela, R., 2002, 'The We-mode and the I-mode', in F. Schmitt (ed.), *Socializing Metaphysics: The Nature of Social Reality*, Rowman and Littlefield, Lanham, Md, forthcoming
- Tuomela, R. and Balzer, W., 1999, 'Collective Acceptance and Collective Social Notions', *Synthese* 117, 175-205