

WHAT IS COOPERATION?*

1. INTRODUCTION

We can cooperate to build a house or to write a book. In such cases cooperation is a joint action. In some other cases cooperation can fall short of being a proper joint action; e.g., the drivers in the street cooperate by following traffic rules, but they do not always act jointly (at least in the sense in which I will be speaking of joint action). While some comments on such cooperation will be made later in the paper, the focus of this paper will be on cooperative joint action. This latter notion will first require a discussion of joint action and joint intention. After this discussion, a kind of “two-dimensional” account of cooperation will be presented. According to it, every joint action type – be it cooperative or noncooperative – can be performed both cooperatively and noncooperatively. A full-blown singular cooperative action is a cooperatively performed token of a cooperative action type.

2. JOINT ACTION

Consider a situation in which some agents are performing a joint action, e.g., to carry a table upstairs, sing a duet, build a house, or get married. How are joint actions built out of individual actions? As a first shot, we can take a joint social action in its broadest sense to be an action performable by several agents who share a “we-attitude” (involving a joint goal, belief, or the like) and acting on this we-attitude (cf. Tuomela, 1992). In this paper we will require in analogy with the single-agent case of a joint action, furthermore, that it be based on joint intention (basically a shared “we-intention” about which there is mutual belief). Some agents’ jointly carrying a heavy object, singing a duet, or toasting someone are typically examples of joint action. In contrast, there is “coaction”, collective action in which agents without having a joint intention have the same goal, perhaps mutually believing so and possibly interacting in various ways. For instance, a proper social norm may require people on a festive occasion to stand up at some

point. This is normally not a joint action but only a coercion. Each person may intentionally stand up because of the norm, but there need not be a joint aim or intention involved. Intentional joint action requires such a joint intention.

It can be argued that a joint action, *X*, and the joint intention to perform it in the “proper” or fullest sense of these notions require that the participants have explicitly or implicitly agreed to perform *X* or an action *Y* such that – normally by practical entailment – also *X* becomes involved; or – to cover some exceptional cases – the participants must at any rate believe that they made such an agreement to act jointly. Without agreement-making, the interdependence and the commitments involved in such joint actions and joint intentions cannot be fully captured. This view has been argued for in Tuomela (1992). However, it should also be noted that there is collective (or “joint”) action without agreement-making; and one can have an intention to perform an action jointly with some other persons without having made even an implicit agreement about it. Consider an example of cooperative collective action. Some persons, seeing a (small) bus starting to slide down a slippery hill, together start pushing it up the hill. I might be passing by and might come to form the intention to push the bus together with the others, viz, have an intention with a kind of joint action as its content, and take part in the pushing without any agreement-making with the others. There may be a mutual belief among all the pushers that I also am taking part in the pushing. This kind of joint activity should be regarded as a case of coercion (or a very weak kind of joint action) rather than “proper” joint action (see Tuomela, 1992, Chapter 2 for discussion).

3. JOINT INTENTION

As is generally accepted, a (mere) personal intention involves making up one’s mind. Similarly, joint intentions such as we-intentions (to be discussed below) involve the participants’ having jointly resolved or made up their minds or made a joint plan concerning what to do jointly. As was remarked above, “proper” joint actions are based on (at least mutually believed) agreement, either explicit or implicit. It should be obvious that if we-intentions are to have proper joint actions as their “satisfiers”, they must also be based on agreement-making. But, as

said, we will accept below that joint intentions can concern also joint actions in a wider sense.

In joint-intention formation each agent accepts for himself: "I ought to participate in our doing X together". This acceptance here means that the agent (at least dimly) recognizes the existence of a joint plan to perform X and accordingly commits himself to performing X together with the others. A joint intention on conceptual grounds leads to each agent's acceptance of "I will participate in, or contribute to, our doing X ", based on his acceptance of "We will do X " (the standard expression for joint intentions or "group-intentions", viz, we-intentions and standing group-intentions, as argued in Tuomela, 1992).

Let us consider the central notion of we-intention in some more detail. We-intentions are action-generating joint intentions that agents have in situations of joint action, e.g., when they jointly intend to carry a table jointly. The content of a we-intention can be taken to be something like "to do X jointly" or "we to do X jointly". A we-intention involves the intention to perform one's part of the joint action. We can say roughly that a member A_i of a collective G ("we" for A_i) we-intends to do X if and only if A_i (i) intends to do his part of X (as his part of X), (ii) has a belief to the effect that the 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 – or at least among those participants who do their parts of X intentionally as their parts of X – to the effect that the joint action opportunities for an intentional performance of X will obtain.¹

Next consider a schema of practical reasoning that a we-intending agent is required to satisfy. This schema also serves to account for the commitments the participants of a joint action have towards other participants (cf. Tuomela, 1984, Chapter 2, Tuomela and Miller, 1988, and Tuomela, 1992, Chapter 3):

- (W)(i) We will do X .
- (ii) X cannot be performed by us unless we perform action Z (for instance, in the case of a cooperative action type X , teach agent A , who is one of us, to do something related to his performance of actions required of him for X).
- (iii) We will do Z .

- (iv) Unless I perform *Y* we cannot perform *Z*.
- (v) I will do *Y* (as my contribution to *Z*).

This schema, to be imposed on all we-intentions and to be exhibited by the we-intenders' dispositions to reason in appropriate circumstances, expresses part of what is involved in saying that a we-intention involves a joint commitment to act so as to contribute towards the realization of the content of the we-intention (viz, "we will do *X*" or, better, "we will do *X* in accordance and because our shared we-intention to do *X*"). Basically, we are jointly committed to perform extra actions required for our joint action *X*, where the extra actions (e.g., actions helping to bring about the other participants' part-performances) are actions not foreseen in our joint plan to perform *X* and hence were not included in our preassigned part-actions. Thus the agents' parts *ex post actu* may be different from their parts *ex ante actu*. If the need arises, extra activity for *X* must be undertaken, and clause (iii) – the intermediate conclusion of (*W*) – deals with specifying our extra joint commitment. Clauses (iv) and (v) deal with what – if anything (free-riderism versus fair share!) – I believe and intend to do in this situation requiring extra action. (For that extra contribution of mine to really be a part of *Z* and an *ex post actu* part of *X*, it must in fact be mutually believed to be a part accepted by the participants.)

Joint actions may be divided into (fully) cooperative and noncooperative actions. In the former type of action, in contrast to the latter, it is in accordance with the participants' preferences to help (assist) other members in their performances of their parts. As seen, in any joint action *X* the participants are jointly committed to furthering *X* and to seeing to it that it comes about; they are thus also committed to performing those extra activities concerning their fellow participants' part-performances on which the success of *X* in their view depends.² While (*W*) indeed applies to all cases of joint action, in the case of cooperative action types *X, Y* may be a helping action in the sense that it helps a fellow participant *A* to perform his part well. (The extra action *Y* may concern the bringing about and maintaining the external circumstances needed for the part-performance itself, or it may itself become a part of *A*'s *X*-related part-action, this latter part-action becoming itself performed as a joint action.) In the case of fully noncooperative action types *X*, viz., actions involving opposing preferences, *Z* and *Y* can not, as long as the agents act on their preferences, be or involve helping *A*

to succeed in his part-performance, but can at best be related to his being able to perform his part at all (so that the joint action X becomes possible). Both in the case of cooperative and of noncooperative action types, the particular behaviors involved (performances of preassigned parts as well as of the extras) can be performed either cooperatively or noncooperatively. In the case of cooperatively performed joint actions X (be X a cooperative or a noncooperative type), the agents perform their preassigned parts and extras with a cooperative attitude and – in addition to performing extra actions required for X – are assumed to be disposed to perform unrequired, cost-involving extra actions (at least extras whose cost do not exceed the gains accruing from them). In the case of noncooperatively performed actions the participants perform their parts and extra actions (only required extra actions can be involved) with a noncooperative attitude, trying to free-ride and minimize their contribution to X .

Of course, in many cases – when the world “cooperates” and things go as expected – there will not be any such extra required actions Z and Y in the sense of our schema (W). The schema anyhow shows what kind of practical inferences the we-intending agents should be disposed to make. In the case of cooperative actions the we-intentions satisfying (W) in the specified way accordingly can be called cooperative.

4. AN ANALYSIS OF COOPERATIVE JOINT ACTION

As seen, joint action types can be cooperative or noncooperative. The performance of a fully cooperative joint action can be taken under favorable conditions to give the participating agents (at least in the case of an optimal choice of agents) jointly a better result – reward or utility – than they can attain by acting separately. In addition to this “input” condition there is the motivational “output” condition – a rationality condition of motivation – that the resulting joint utility can be expected to be divided among the participants so that nobody loses when compared with the situation of acting alone – in those cases in which the joint action indeed gives higher joint utility than the sum of the utilities accruing from separate action. (There are plenty of cases where this output condition is not satisfied in actual life; cf. cases when the participants have been hired or coerced to perform a part.) Note that when

a collective good is being produced everyone in the group can enjoy the result of the joint action without limit, in principle.

The “better result” spoken about in the input condition may not come about in every normal circumstance, because of a lack of ability or of the right opportunity. Nevertheless, this feature indicates the central part of what the main reward is in a cooperative joint action. What is even more important is that in the instances of a fully cooperative action, help – in the sense of actions strictly contributing to other participants’ performing their parts well – is always in accordance with the participants’ inbuilt preferences, and, at least in some cases, the more the participants actually help each other, the more successful the joint action will be, other things being equal. (We must assume in the case of rational agents that the costs of helping actions do not exceed the gains accruing from them; given that X indeed will be performed, such cost calculations will be nontrivially involved only in the case of unrequired extra action.) We can say that in the case of cooperative action we (tend to) “stand or fall together” – the participants’ preferences (*qua* participants) are highly positively correlated; and this explains the possibility of helping the other participants (disregarding considerations of the cost of helping actions). Carrying a table upstairs and building a house are examples of fully cooperative actions.

In a noncooperative action type, in contrast, there is only the joint action “bottom” – in the case of proper joint action an agreement to perform the action. However, beyond that (in a different sense “cooperative”) bottom there is no way to help other participants with their part performances (although it is possible to help to bring about and maintain “joint action opportunities”). Competitive games such as chess or organized fights are examples of this kind of noncooperative joint action. In them a participant can at best – and, indeed, ought to, when needed – contribute to the preconditions of the other participants’ part-performance, but cannot help them improve their performance, given that the agents act on their part-related preferences (utilities) built into the structure of the action. Another subtype of noncooperative action is joint action based on exchange – cf. selling and buying, where performing the exchange can be taken to be the agreement-based “bottom action”. Joint actions can be performed with various underlying attitudes. Thus any joint action, no matter whether cooperative or noncooperative, can – on conceptual grounds – be performed cooperatively or noncooperatively.

On the basis of this discussion we can present the following classification. First *joint action types* (or contexts or structures) can be *cooperative* or *noncooperative*. Let us list the ingredients that we have discerned in it earlier in this paper. First there is a participant's (subjectively and objectively required) *part-action* (assume this to be sufficiently clear for the purposes of this paper, cf. Tuomela, 1992, Chapter 2). Next there are the *extra actions* that a joint action, X , may require. In the context of schema (W) we spoke of an action Z that is required by X . The performance of Z has to be divided among the participants. Let us call Z and its parts – as divided among the participants – *required extra actions*. There may, in addition, be actions which contribute to X but are not necessary for its successful performance. Let us call these actions *unrequired extra actions*.

Which of these ingredients are present in cooperative actions and which in noncooperative ones? In the case of both cooperative and noncooperative joint action types, part-actions must be involved, and also extra actions of both kinds may be present (depending on the agent-external circumstances of action and, in the case of unrequired extras, also on their relative costs). In the case of fully cooperative action contribution to other participants' performances of their parts is in accordance with a participant's preferences (utilities) built into the joint action and principle of acting on one's preferences, given that the costs related to those helping actions are smaller than the gains expected to accrue from them. Or we can say that disregarding the costs related to helping actions it is promotive of the satisfaction of anyone's part-related preferences in question to help anyone. Note that required extra actions also are included in a successful performance of X , although there may be a dispute among the participants about how to divide Z into parts. As a joint action is at stake, the participants have the collective responsibility to perform what the joint action in question requires. The notion of helping in the case of cooperative action concerns helping other participants not only to perform their parts but to perform them well.

In general the following helping activities may be involved in the case of a cooperative action type: i) contributing to the coming into existence of a precondition of another person's part, ii) contributing to or participating in the performance of the part itself (in the latter case the other person's part is actually performed as a joint action), iii) contributing by counteracting negative interference related to the

other person's part-performance. A fully cooperative joint action allows for both required and unrequired extra actions; in the case of rational action it holds that in the case of the unrequired extras the costs related to them are smaller than the expected gains generated by their performance. (In the case of required helping actions there will be no room for cost calculations, at least as long as their costs do not exceed the specific gain due to the joint performance of X .) In a noncooperative action type extra actions only of kinds i) and iii) can occur.

The central point about fully cooperative actions is that one can help others in their performances of their parts well and make other contributions in the sense that one can contribute positively to the joint utility. In a cooperative joint action everybody is assumed to contribute positively to the joint utility, which is divided among the participants if not evenly at least by a monotonously increasing function of the part-performances (or, better, the utilities of part-performances). We recall also that, under favorable condition, each participant can expect to benefit from any other participant's contribution to X . In game-theoretic terminology, games with coinciding or nearly coinciding interests roughly correspond to this notion of a cooperative action type, given that there is an agreement to play the game.³

What about noncooperative joint action types, viz. actions involving more or less opposing interests? In their case – *qua* their being joint action types – the participants must perform their (preassigned) parts plus the required extra actions (contributive actions related to the bringing about and maintaining of joint action opportunities), and they may also perform some unrequired extra actions. The central difference between cooperative and noncooperative actions is that in the case of the latter it is not at least always prudentially possible for a participant to help the others with their performances of their parts well or with anything that specifically contributes to their part-performances. That is, this is not possible relative to such a participant's preferences (utilities) and the assumption that a participant *qua* a participant acts on his preferences. Competitive joint actions and exchange-actions are examples of noncooperative actions. In a noncooperative action the preferences (or utilities) of the participants (part-performers) are at least to some extent antagonistic: helping another participant will lower one's gain (payoff, utility) in the joint action. In game theory, zero-sum games and many mixed-interest games belong among noncooperative actions taken in the present sense. (One may speak of noncooperation

also in many other senses, e.g. in the sense of blocking or of negative interference; we will not here consider these ideas.)

Next we distinguish between *cooperatively* and *noncooperatively* performed joint actions (viz., action tokens). All joint actions (thus both cooperative and noncooperative action types) can be performed either cooperatively (viz., with a cooperative action-disposition or attitude) or noncooperatively. A person having a cooperative attitude towards a joint action X must be disposed to reason and act in ways contributive to X . He must thus be disposed willingly to perform his part of X and must also willingly accept a share – reasonable for him relative to his capacities and skills – of the required extra action Z ; and he must be willing to perform unrequired extra actions related to X (as long as their costs do not exceed the accruing gains). In contrast, a person with a noncooperative attitude towards X is disposed to free-ride and to minimize his contributions to X . He will reluctantly perform his part of X , at least if his part-performance indeed is necessary for the coming about of X . Otherwise he could not be a participant in X at all. He will be reluctant to accept any required and unrequired extra tasks. (Often an “Italian strike” belonging to a broader joint action type, usually a cooperative one, is performed with a noncooperative attitude.) The presence of a cooperative attitude – as contrasted with the presence of a noncooperative attitude – clearly may affect the effectiveness and speed of performance and the general “social atmosphere” and amount of “we-feeling” in the group.

On the basis of our discussion, the following somewhat more exact characterizations of the “pure” notions of a cooperative action type and a cooperatively performed singular action can be proposed:

- (CAT) X is a *cooperative joint action type* if and only if for some number m (≥ 2) and some choice of m agents,
- (1) X is a joint action type;
 - (2) on conceptual-normative and/or factual nomological grounds, those m agents can jointly perform X more rewardingly (viz., as giving the participants no less utility) under favorable conditions (conditions related to an optimal performance of X , irrespective of the participants’ utility gains or losses not related to the type X) than when X is performed separately by them, provided that X can be performed separately at all;

- (3) disregarding the costs of helping actions, it is at least under favorable conditions promotive of the satisfaction of any participant's part-related preferences (utilities) and the principle of acting on one's preferences, to help any participant to perform his part (indeed, to perform it well) and to perform extra helping actions.
- (C) x is a *cooperatively performed* intentional joint action token if and only if
- (a) x is a token of a (cooperative or noncooperative) joint action type X ;
- (b) x is performed intentionally by the participants (on a relevant we-intention);
- (c) x is performed intentionally with (and partly because of) a cooperative attitude, viz., in our schema (W) it is assumed that the participants perform actions regarded as required for X – be they preassigned part-actions or other required actions – and perform them willingly; and the participants are also assumed to be disposed willingly to perform unrequired but contributive actions, thus being disposed to take extra costs (this being rational as long as the costs from performing them are smaller than the gains accruing from their performance).

A few clarifying remarks are in order. The notion of joint action type used here need not be regarded as ontically significant. It merely signifies the idea that there are recurrently performable actions X which may involve certain preference structures. As to clause (1) of (*CAT*), the notion of joint action is to be regarded basically as one performable on a we-intention. This clause does not do much work in the present analysis.⁴ Concerning clause (2), it might be thought that it primarily relates to X being a joint action type rather than to its being specifically a cooperative joint action type. Such noncooperative joint actions as tennis-playing could satisfy it on conceptual grounds, while many other actions such as carrying a heavy table satisfy it on factual grounds. This clause comes to say that the cooperative joint action type X requires several participants for an optimal performance of it. Note that it is not required by (2) that the increase of the number of participants makes X somehow more successfully performable – that requirement is, in a sense, built into clause (3). Actually (2) is simplified. For one

thing, it does not explicitly take into account the part-division of X into parts. However, it is compatible with the existence of a preassigned task-division. The notion of a favorable circumstances means that the performance of X in such circumstances is expected to give an optimal performance of X . For instance, this requirement guards against the occurrence of a "crowding effect", lowering the joint utility, with more than m participants. We assume here that there are stable part-related preferences belonging to X . The utility related to X concerns them rather than some idiosyncratic, personal action-preferences of the participants. The X -related favorable conditions must nevertheless be taken to include considerations related to the participants' abilities and to the external physical and social "joint action opportunities". Understood along these lines the notion of favorable circumstances becomes informative. Note that not all joint actions then satisfy clause 2). Think, for instance, of norm-based institutional actions that the participants even *qua* participants prefer to perform separately if at all. Such joint action would possibly be less rewarding, not only psychologically and socially but also with respect to the goodness of result, than when performed separately by those agents or some subset of them. For instance, some versions of democratic decision making seem to qualify as examples of such norm-governed institutional actions.

Clause (3) says that, at least under favorable conditions, help is prudentially possible in the case of cooperative actions. Consider the proverbial case of the many cooks who jointly spoil the broth. We can take the spoiling of the broth to result from the lack of joint skills of the cooks: they do not cooperate sufficiently skillfully. But if they have that joint skill and if also the other favorable conditions obtain, they would jointly be better off (or at least not worse off) than each acting separately, clause (2) becoming fulfilled. It is also possible for them to help each other in cooking, and thus clause (3) is seen to be fulfilled. Indeed, the latter fact explains in part the fulfillment of clause (2). Thus, m -person cooking is a cooperative joint action type.

As indicated, the possibility of helping in the context of a cooperative joint action type means that every participant's part-performance or an 'extra' action (related to the joint action X) can be assumed to contribute positively to the joint utility, which is divided among the participants if not evenly at least by a monotonously increasing function of the part-performances (or, better, utilities due to the part-performances). The utilities are "positional" or "part-related" (rather than personal). Nat-

urally, the participants may have personal preferences not related to the action X but which nevertheless X can help them satisfy. Such personal preferences may be in partial conflict with each other.

Note that there can also be individual action in a cooperative situation with a common goal, this action being based on mere mutual belief. However, this kind of cooperative coaction – resembling cooperative action as described by (2) and (3) of (*CAT*) – falls short of proper, agreement-based joint action. For instance, agents said to mutually “cooperate” rather than “defect” in a situation of social dilemma such as exemplified by the Prisoner’s Dilemma or the Chicken game are often taking part only in interactive coaction. In a two-person Prisoner’s Dilemma mutual cooperation means mutual choice of the cooperation-alternative (C), which results in the collectively Pareto-preferable joint outcome CC . Such cooperative behavior can be based on a cooperative attitude (readiness) to accept collectively beneficial goals (such as CC) for one’s group or on a “sociality-readiness” so to act together and in the company of others. The other possibility is, of course, that they “cooperate” due to an agreement that they have made. The problem with agreement-making here is that it is hard to achieve binding agreements, as each participant has an incentive not to participate in the agreement, as soon as he comes to believe that the others will contribute (or that, anyway, sufficiently many of them will contribute so that the action X will become performed).

The analysis (C) seems not to require further remarks in this connection. In conclusion, our two-dimensional account of cooperation has been summarized by the analyses (*CAT*) and (C), representing the two dimensions in question. We may say that the two dimensions of a cooperative joint action type X are (1) the possibility to help other participants with their performances of their parts of X and (2) the performance of one’s part of X with a cooperative attitude. (We also referred to an additional “cooperative” aspect: the participation in the performance of the actions that the joint action requires and in the case of proper joint action especially the fulfillment of the agreement to perform X .) A joint action token x of an action X is fully cooperative if (i) X is a cooperative joint action type and (ii) x is performed with, and in part because of, a cooperative attitude. (Clause (ii) involves the performance of helping-actions whenever needed.) Such a token can be only *weakly* cooperative in two quite different senses: either x is a noncooperatively performed token of a cooperative joint action type

X , or x is a cooperatively performed token of a noncooperative joint action type X (or, better, a joint action type X which is not cooperative).

A fully cooperative joint action can be performed under coercion. A person can coerce somebody else to perform an action jointly with him, e.g., into a joint business venture, once the sanction is sufficiently severe. The coerced person may willingly participate and save himself from such severe punishment. (It does not matter whether the coercer is one of the participants of joint action or some outsider.) This means that the notion of a cooperative attitude has nothing to do with altruism or related psychological attitudes.

Noncooperative joint action types can be technically characterized largely by analogy, and they will include not only actions with opposing preferences (and thus “zero-sum” situations) but also many kinds of actions with mixed inbuilt preferences or interests. Recall that in the case of purely noncooperative action types it is essential that it is not possible to help other participants’ performances beyond what the mere coming about of the joint action requires. Noncooperatively performed joint action (be it of a cooperative or of a noncooperative type) can obviously be characterized by analogy with (C), negating the requirements of the presence of a helping attitude and unrequired actions in its clause (c).

5. CONCLUSION

In this paper a kind of “two-dimensional” account of cooperation has been presented and argued. This analysis builds on two distinctions: the distinction between cooperative and noncooperative action types and the distinction between cooperatively and noncooperatively performed joint action (be such an action an instance of a cooperative or a noncooperative joint action type).

A fully cooperative joint action type is one allowing that the participants help each other to perform their parts well, whereas that is not prudentially possible in the case of noncooperative action types. A cooperatively performed action token is one based on a cooperative attitude – willingness to perform what the joint action in question requires and in addition also unrequired contributive actions.⁵

NOTES

* I wish to thank Kaarlo Miller and Maj Bonnevier-Tuomela for helpful discussions. Criticisms by the members of the Department of Philosophy, Logic and Scientific Method at the London School of Economics and Political Science when presenting this paper at the departmental research seminar were helpful, too.

¹ The reader can be referred to my 1992 book as well as to Tuomela and Miller (1988) and Tuomela (1991) for a discussion of the notion of we-intention and the notions involved in the above characterization. Here the analysis is mentioned only to illustrate how to make the agreement view more detailed. Obviously shared we-intentions can be called joint intentions. In Tuomela (1991) and (1992) I discuss also a somewhat wider notion of joint intention, viz., shared "group-intentions". A group-intention is an intention applicable to a single agent and expressible by a conative sentence "We will do X ". The agent's having the group-intention to do X amounts to his ("conatively") endorsing "We will do X ".

² As is generally accepted in the literature, ordinary intentions are commitments to action. Similarly, we-intentions are the participants' commitments to joint action. In the case of proper joint actions and the corresponding joint intentions we have joint commitments in a full sense. As our focus will be on proper, agreement-based joint (and cooperative) actions and we-intentions to perform them, we can say a little more about this central case. The notion of an effective agreement involves the participants in the agreement being jointly committed to satisfy the content of the agreement, here performing the joint action X . Such a joint commitment involves that the participants (1) are committed to perform their parts of the joint action, (2) are responsible to each other for performing their parts, and (3) are committed to furthering X if they believe their extra contributions (e.g., help, pressure) are needed. We may also loosely say that each we-intending agent must have as his purpose that X be performed, but – in order not to make our analysis viciously circular – here this will only be understood as sharing the joint commitment towards performing X . It is worth emphasizing again that commitments (1), (2), and (3) are based on the interpersonal obligations ensuing from the underlying obligation, and thus are stronger than the commitments related to mere contentions, shared intentions towards X , about which there is mutual belief. Thus also, e.g., the commitment to act as schema (W) specifies need not be satisfied in the latter case: only a collection of separate personal commitments (which may be changed merely on personal grounds) is entailed. This gives support to our thesis about the existence of an underlying agreement. The "extra" actions Y of schema (W) clearly indicate A_i 's commitment to X and the accompanying task to follow through the performance of the joint action until its end.

³ Harsanyi, 1977, p. 109, characterizes a game with strictly identical interests as one in which "each player's payoff is a strictly increasing function of any other payoff over the entire payoff space P of the game".

It may be noted that in game-theory the distinction between cooperative and noncooperative games is quite different than in this paper. Game-theorists call a game cooperative if the players are allowed to communicate and make agreements; otherwise it is noncooperative. But both cooperative and noncooperative proper joint actions in our sense are always based on agreement-making, and are thus cooperative in the game-theorist's terminology.

A somewhat broader notion of preference-similarity or preference correlation is the

following probabilistic account, which I regard as acceptable for all cooperative and noncooperative actions. We can explicate the conditions underlying the possibility of helping in terms of conditional probabilities of goal-attainment. Consider the case of two agents, A and B , performing an action X jointly. We assume that G_i , $i = a, b$ being the X -related payoff i will get when X is being performed (here it may or may not be necessary that each agent perform his part of X for X to become performed). In the case of pure cooperative action situations we can impose the following simple probabilistic clauses, in which X_i means performing one's part well and where $\neg X_i$ means performing one's part only at a minimal, required level for X to come about:

- (i) $p(G_a/X_a \& X_b) \geq p(G_a/X_a \& \neg X_b)$
- (ii) $p(G_b/X_a \& X_b) \geq p(G_b/\neg X_a \& X_b)$

In the case of noncooperative action types the inequality signs will be reversed, \leq replacing \geq .

⁴ Accordingly, the notion of joint action type can here be understood in the we-attitude-based sense of (5.8) of Tuomela (1984); also cf. Chapter 2 of Tuomela (1992).

⁵ The only conceptual clarification of the notion of cooperation I am aware of is given the recent paper by Bratman (1991). At least from the point of view of the present theory of cooperation Bratman's account is open to several criticisms. See Tuomela (1993).

REFERENCES

- Bratman, M. E.: 1991, 'Shared Cooperative Activity', *The Philosophical Review* **101**, 327–341, forthcoming.
- Harsanyi, J.: 1977, *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, Cambridge University Press, Cambridge, Mass.
- Tuomela, R.: 1984, *A Theory of Social Action*, Synthese Library, Reidel Publishing Company, Dordrecht and Boston.
- Tuomela, R.: 1991, 'We Will Do It: An Analysis of Group-Intentions', *Philosophy and Phenomenological Research* **LI**, 249–277.
- Tuomela, R.: 1992, *The Importance of Us: A Philosophical Study of Basic Social Notions*, book manuscript accepted for publication by Stanford UP.
- Tuomela, R.: 1993, 'Cooperation and Joint Action', manuscript.
- Tuomela, R. and Miller, K.: 1988, 'We-Intentions', *Philosophical Studies* **53**, 115–137.

Manuscript submitted April 6, 1992

Final version received August 18, 1992

Dept. of Philosophy
University of Helsinki
Unioninkatu 40B
00170 Helsinki, Finland