

RATIONALIZING COORDINATION  
TOWARDS A STRONGER CONCEPTION OF COLLECTIVE INTENTIONALITY

Hans Bernhard Schmid  
University of St. Gallen  
Hans\_Bernhard.Schmid@unisg.ch<sup>1</sup>

With very few exceptions only, philosophy of mind has exclusively been concerned with *individual* intentions. Not all intentionality, however, is of the form *ego cogito*. There are intentional states of the form *nos cogitamus*, too. Among the most exiting recent developments in the philosophy of mind is the turn to a systematic analysis of the structure and role of *collective intentions and beliefs*. Philosophers such as Raimo Tuomela (1988; 1995), Margaret Gilbert (1992), John Searle (1995), and Michael Bratman (1999) have put the analysis of shared intentionality on the map of philosophical thinking. Even though no consensus concerning the structure of collective intentionality has emerged as yet, this movement has substantially widened our view of the mind.

The aim of this paper is twofold. Firstly, an especially conspicuous example for the importance of this development to the social sciences in general and economic theory in particular shall be given, i.e. the theory of *coordination* (sect. i-iii). Intuitively speaking, coordination problems differs from *cooperation* problems in that intuitively speaking, rational participants aim at “matching” their individual choices among the available alternatives, so that there is no incentive for unilateral defection. Whereas the structure of cooperation (as exemplified by the *Prisoners’ Dilemma*) has been at the very center of thorough philosophical analysis – as well as of experimental work in economics – over the last half of a century, the structure of *coordination* has not received nearly as much attention. One reason for this is that in most real life cases, coordination problems are easily solved by means of *conventions*. In line with some recent literature on the topic, I will claim that for its individualistic limitations, the standard economic model of human behavior fails to explain how conventions make coordination among rational agents possible.

---

<sup>1</sup> This paper was originally presented at the 12<sup>th</sup> International Congress of Logic, Methodology, and Philosophy of Science in Oviedo, Spain, on August 11, 2003. Parts of the argument developed here first appeared in Schmid 2003a. Some passages of section v. of this paper are taken from Schmid 2003b. I am very grateful to Barbara Montero and Mark D. White for sharp comments earlier versions of this paper.

In the second part of the paper (sect. iv-vi), it is argued that even though the existing accounts of collective intentionality point the way towards an adequate account of coordination, a more robust conception of collective intentionality than the ones to be found in the existing literature is needed. In a discussion of Robert Sugden's theory of *Team Thinking*, Michael E. Bratman's account of *shared cooperative activity*, and John R. Searle's theory of collective intentionality, three features of an adequate account of collective intentionality are introduced. The paper concludes by addressing an objection that is often raised against more robust conceptions of collective intentionality, i.e. the problem of the group mind.

### *I. The Trouble with Deconditionalization*

For all of its blatant absurdity, the following fictional story might serve as an introduction to the problem. On a sunny afternoon, the police are called to the site of an accident. A look on the scene immediately makes clear what has happened. On a curvy street, a car has come into the oncoming traffic's way, resulting in a frontal collision. Luckily, no one is seriously hurt. The police question the culpable driver. Why did he come into the other driver's way? Did he lose control of his car, did he experience any technical difficulties prior to the accident, was he inattentive or somehow distracted? The driver answers in the negative. He had, he declares, become aware of the oncoming car early enough, he was in full control of his car at all times, he knows the traffic rules that were in no way violated by the other driver. He has no suicidal or otherwise destructive impulses or desires, nor any reason to assume the other drivers had some such preferences. So how come he ended up on the wrong side of the street? He answers with a stern expression on his face: "I just couldn't see why I should keep to my own side of the street rather than swerve to the other's in order to avoid a collision."

For a moment, the police officer in charge loses his straight face in disbelief, and a disparaging remark concerning the driver's state of mind slips his mouth. Now it is our driver's turn to lose his calm. Angrily and firmly he asserts, that he knows perfectly well that he is the guilty party in terms of the law, and that he will accept all charges in terms of legal accountability, but that he rejects any accusations of irrationality. He explains that in the hindsight, he knows very well that he had better chosen to stay on the right side of the street. But that does not mean, he claims, that his decision to swerve to the other side was *irrational* at the time it was made. For when the oncoming car approached, he had to decide over the two strategies "right" and "left". As opposed to "right", "left" is illegal and thus connected with the risk of getting fined by the police. At the same time, however, it would be plainly

irrational *not* to commit a minor traffic violation such as choosing “left” if it is the only way to avoid an accident. “Thus I knew”, our driver concludes, “that if the other driver had decided to swerve to the left, it would have been rational for me to swerve to the left, too. And I assumed that since all of this is common knowledge, and since the other driver is a rational being, too, he must be having the same thoughts, and thinking that I have them too, and so on and so forth. It dawned on me that however deep my (and her) analysis of the situation would dig, what was the *rational* thing to do in our respective shoes would always remain *hypothetical*. In such situations, you can say what’s rational only *if* you have some expectation concerning the other’s decision. But at the same time, you know that the other’s decision in turn depends on his expectation concerning your own choice. Thus in my reasoning, I got stuck in a circle of ‘ifs’, and when we finally reached the point where each of us had to make her final decision over our strategies, I saw there was no way to derive from all that what was *in fact* rational for the other driver and for me to do. You can’t say the other driver’s choosing ‘right’ was rational and my choosing ‘left’ irrational. For as it turned out, her decision to stick to the rules turned out to be wrong, too, since she could easily have avoided getting stuck in this car accident by swerving to the left!”

Of course there is something wrong with this claim. This whole way of reasoning is profoundly mistaken. From a commonsensical understanding of rationality at least, we should take the police officer’s side for once and admit that he is probably right when he doubts our driver’s rationality. To the pre-theoretical eye, it is obvious that it is *plainly irrational* not to keep to the right where the right-side rule applies, if one assumes that the rules as well as the absence of suicidal preferences are common knowledge. But is this strong (and presumably reasonable) pre-theoretical intuition backed by our standard theory of rationality in action, such as implied in the economic model of behavior? Can our driver be proven wrong within the theoretical framework of individual expected utility maximization?

Our driver is, of course, well aware that there are not only two coordination equilibria in the game at hand (“right”/“right” and “left”/“left”), but also a *convention*. In David Lewis’ (1969) terms, “right”/“right” is *salient*, or, in Thomas C. Schelling’s (1960) term, a *focal point*, which makes, in a sense, “right” the *obvious choice*. Our driver does not ignore the existence of the right side traffic rule or deny that this rule is common knowledge among the participants. Also, he might easily grant that it is clear from previous experience that the probability of the strategy “right” being chosen by the drivers in a given population is extremely high. What he is getting at is that mere objective behavioral regularities as such and the existence of precedents do not provide a *reason* for a rational choice. Our driver’s point is

that he could not form any *prior expectation* concerning the other's behavior because he knew that as a *rational being*, the other driver would have to base his decision on his expectation concerning our driver's behavior, and not just blindly follow some behavioral pattern. And indeed it's hard to see why it should be rational not to treat the other not as a rational being, but as some sort of compulsive salience seeker (we will come back to this below). If, however, the other party is treated as a rational being, it is clear that it is rational for her to conform to the general pattern *only if* she expects that the same will hold true for our driver himself, which sets off the infinite regress of interdependent expectations.<sup>2</sup> In Raimo Tuomela's words, what our driver is getting at is the "deconditionalization problem" in the theory of coordination (Tuomela 2002: 388). "Deconditionalization" here means: getting rid of the *condition* deriving on which "right" is the rational choice, i.e. the big *if*: "right" is the rational choice *if* one has no reason to expect the other to choose "left". It seems impossible, however, to derive from the *hypothetical* (or: "conditional") rationality of "right" that given common knowledge of the right side traffic rule and the absence of suicidal preferences, "right" is in fact the rational choice.<sup>3</sup>

Deconditionalization can be a *real life problem* in situations where there are no conventions, and no salient solution or focal points (as in the case of two pedestrians colliding on the sidewalk because of a mismatch of the chosen strategies). Where there are commonly known conventions (such as in the case of motorized traffic), however, these problems do not empirically occur. It seems that in these latter examples, deconditionalization is easily achieved. We do not experience serious difficulties in deciding whether or not to stick to the

---

<sup>2</sup> For a very clear formulation of the problem cf. Parsons/Shils 1959: 105: "In interaction ego and alter are each objects of orientation for the other. The basic differences from orientation to nonsocial objects are two. First, since the outcome of ego's action (e.g. success in the attainment of a goal) is contingent on alter's reaction to what ego does, ego becomes oriented not only to alter's probable overt behavior but also to what ego interprets to be alter's expectations relative to ego's behavior, since ego expects that alter's expectations will influence alter's behavior. Second, in an integrated system, this orientation to the expectations of the other is reciprocal or complementary." It is very telling that in Talcott Parsons's view, this "double contingency" implied in interaction cannot be adequately dealt with within the action theoretic framework of the standard economic model of behavior, even though one might be sceptical as to whether Parsons' own systems theoretical framework is more successful in explaining the structure of coordination.

<sup>3</sup> "An agent cannot rationally (...) form and satisfy his action intention without a circular reference to the other agent's intention. Hence, he cannot, so to speak, finitarily infer or compute the satisfaction value of statements like 'I will do X if you will do X' in the kind of coordination situation on the basis of independently assignable satisfaction values of 'I will do X' and 'you will do X', because there simply are no such satisfaction values" (Tuomela 2002: 390).

traffic rules in order to avoid a collision, where the rules and the absence of suicidal preferences are common knowledge. Our pretheoretical intuition is that in a very basic sense, this is simply a matter of *rationality* in terms of straight reasoning. One of the philosophical questions behind this is the following: How precisely do *salience* and the existence of *conventions* provide us with reasons? Why is it rational to choose the salient strategy in pure coordination games, when apparently, the existence of conventions, of focal points, or of salient solutions do not *per se* solve the deconditionalization problem? Or, put negatively: What precisely is wrong about our drivers claim that right side traffic rule or not, any attempt to base a rational decision on salience is immediately drawn back into the infinite regress of interdependent expectations?

*II. The Principle of Coordination*

In his very influential paper on the topic, David Gauthier (1975) addressed this question. Going beyond Lewis and Schelling, his brilliant move was to draw attention to the role of the *description* under which the players perceive their available strategies.<sup>4</sup> For drivers who do not know any conventions the alternatives at hand are simply “right” and “left”.

The situation might look like this:

	left	right
left	0/0	-1/-1
right	-1/-1	0/0

Fig. 1: pure coordination

For rational players (“rational” in terms of individual expected utility maximization) it is clear that according to the “principle of insufficient reason”, they will choose randomly between the two strategies, which makes all four possible outcomes equally probable. If, however, “right” is *salient*, the players’ choice is now a different one: it is between the alternatives “choose salience” (i.e. choose “right”) and “ignore salience” (i.e. choose randomly), which has a single coordination equilibrium, viz “choose salience”/“choose salience”.

After the “relabeling” of the available strategies, the situation looks as follows:

	choose salience	ignore salience
choose salience	0/0	-0.5/-0.5

---

<sup>4</sup> For a more detailed account of Gauthier’s approach cf. Sugden 1995.

ignore salience	-0.5/-0.5	-0.5/-0.5
-----------------	-----------	-----------

Fig. 2: coordination, re-labeled strategies

Thus by means of salience, a game with two equally good coordination equilibria is miraculously transformed into a game with only one best coordination equilibrium, which makes it rational to choose the according strategy, i.e. stay on one's own side of the street. All of a sudden, the deconditionalizing problem seems to disappear. However, this transformation is too good to be true, and our driver will have no hard time proving Gauthier wrong (cf. Provis 1977; Miller 1991; Gilbert 1996; Goyal/Janssen 1996). For Gauthier simply ignored that in the transformed version of the game, the choice is not between two, but between three strategies. The options open to the players are not just *either* to observe salience (i.e. choose "right") *or* to ignore salience (i.e. choose randomly); the third option is to *choose the non-salient* (i.e. choose "left").<sup>5</sup> In other words: in a derivative sense, the existence of *salience* makes the non-salient strategy salient (one might speak of *secondary* or *derivative* salience). This third strategy, however, has another equally good coordination equilibrium ("choose the non-salient"/"choose the non-salient"), which throws us right back into the initial deconditionalizing problem.

The "re-labeled" situation is the following:

	observe salience	ignore salience	choose the non-salient
observe salience	0/0	-0.5/-0.5	-1/-1
ignore salience	-0.5/-0.5	-0.5/-0.5	-0.5/-0.5
choose the non-salient	-1/-1	-0.5/-0.5	0/0

Fig. 3: coordination, re-labeled, complete set of strategies

Another line of argument in Gauthier's paper goes as follows. The drivers could get by each other either by both choosing "right" or both choosing "left". However, "right"/"right" is not only the *salient* solution. It is also *payoff dominant*, i.e. better for both, for by choosing "left", both drivers run the risk of getting fined for violation of the traffic regulations even if they manage to get by each other collision-free (the risk of getting caught in the act is low; however, it rises to certainty when a collision results, for then, the police will be called in and

---

<sup>5</sup> Provis has a convincing explanation of why this obvious weakness of his transformation argument could have slipped Gauthier's notice. "[The existence of the third option] is obscured because Gauthier introduces his suggested alternatives as being choosing the salient option and ignoring salience. That phraseology diverts attention from the fact that one way of not ignoring salience on an option is by performing the non-salient option *qua* non-salient option" (Provis 1977: 509).

the culpable driver will be fined). Thus the two coordination equilibria in the initial game are not equally good after all, which means that “right” is *weakly dominant*.

Thus a more adequate description might reveal an asymmetry between the two coordination equilibria in the original situation:

	right	left
right	0/0	-1/-2
left	-2/-1	-0.1/-0.1

Fig. 4: coordination, unequal equilibria

As compared to “right”/”right”, “left”/”left” is Pareto inferior. As Gauthier states in his famous “Principle of Coordination”, this makes it rational to choose the according strategy:

„In a situation with one and only one outcome which is (...) a best equilibrium, if each person takes every person to be rational and to share a common conception of the situation, it is rational for each person to perform that action which has the best equilibrium as one of its possible outcomes“ (Gauthier 1975: 201).

However, it is not as obvious as it might first appear that the principle of coordination is a *rational* principle, and indeed it seems that what our driver is getting at is that it is not. In spite of the fact that “left”/”left” is Pareto inferior, it is still an equilibrium, i.e. it is rational to choose “left” *if* one expects the other to choose right, too. Thus it appears that just as in the case with two equally good equilibria, both strategies are *hypothetically* rational. Or, in other words: it is rational for the single participants to adopt the *Principle of Coordination* only *if* she or he expects it to be adopted by the other participants, too. Thus once more, the participants see themselves confronted with the deconditionalizing problem. From the participant’s perspective, the question “why does rationality require *me* to follow this principle?” remains open (Sugden/Hollis 1993: 11).

This, however, did not slip Gauthier’s notice. Against possible objections of this type, Gauthier puts forth the following argument: Because of the coincidence of the individual payoffs in both equilibria, the participating individuals can identify their respective individual choice over the two possible strategies “right” and “left” with a collective choice over the two equilibria. Since “right”/”right” is better for both than “left”/”left”, it seems clear that “right” is the rational individual choice. In a much similar vein, Maarten C. Janssen has argued that “right” is the rational choice since whoever chooses “right” chooses the “better plan” (Janssen 2001; 2001b).

Will that finally convince our driver? There is reason for doubt. That “left”/”left” is Pareto-inferior as compared to “right”/”right” did by no means escape his notice. Thus he will grant that “right”/”right” is the better plan. What he points out, however, is the fact that when the oncoming car approached, he had to make an *individual* decision, not a common decision. It had nothing to do with a decision over plans; it was a decision between strategies. In his view, this difference is a decisive one. For it is all that his reasoning was about: the strategic interdependence of decisions in the given situation. It is simply impossible to derive from the fact that “right”/”right” is the “best plan” in terms of the optimal outcome which would be rationally chosen in a common decision, that “right” is the rational individual decision in the given situation. Whether or not there is a single best equilibrium: as long as there are multiple coordination equilibria, the basic deconditionalizing problem remains the same.

For all that has to be left to a more detailed analysis of the problem at hand, it might have become apparent how stunningly little there is to say against our driver within the framework of the standard theory of rationality in action. If this should turn out to be true, it seems that what we have here is a open conflict between theory and common sense. If the foregoing is right, this tension is not just an *apparent* tension. Either there is something wrong with our pretheoretical intuition, or with the theory.

### *III. The Irrationalist Position*

So who is at fault here – common sense or the standard theory? This is by no means a rhetorical question. Not all of those philosophers who believe that the standard theory of rationality in action cannot be reconciled with the commonsensical intuition concerning the rationality of coordination hold that this is the theory’s fault. Quite to the contrary: some philosophers state that we should let go of our pre-theoretical intuition rather than revising the standard theory. Famously, Thomas C. Schelling seems to be vaguely pointing in this general direction when he states that success in pure coordination games “depend[s] on imagination more than on logic” (Schelling 1960: 57).

It appears, however, that once we start out by saying that focal points or salience do not provide a reason for a rational decision in pure coordination games, we will end up having to grant our driver that he was right after all in saying that there was ultimately no *rational* decision in his particular situation. Such a concession not only ruffles the feathers of our common sense. It also gives rise to the question: how come that with the very peculiar and purely fictional exception of the driver in our example, we reliably manage to coordinate

along the guidelines of the traffic rules, if it is not outright *rational* to stick to those rules? Is this system of conventions on which we rely so much in our everyday dealings really built on mere *imagination*?

Those philosophers who take the *irrationalist* position in the current debate on the structure of coordination usually quote some non-rational *impulse* (e.g. Thalos 1999), some psychological propensities or some “blind” behavior of some other kind (cf. Gilbert 1996). In this view, our driver was by no means *irrational* when he failed to choose “right”; he just happened not to act on the usual kind of impulse.

Is this irrationalist position right? I do not think that we have to go much deeper into the details to encounter some of the difficulties in which this line of argument seems to be running. It is obvious that non-rational impulses (including habits, psychological propensities, or some such) are neither a *sufficient* nor a *necessary* condition for human coordinative behavior. Any student driver who knows the rules, and knows that the rules are common knowledge, will stick to them, even though she has not yet acquired any corresponding habits or extrarational propensities and impulses. It is not the corresponding impulses that make it rational to choose the salient strategy; rather, it is because it is rational to choose the salient strategy that we acquire the corresponding impulses and habits. In other words, the irrationalist position puts the cart before the horse.

This becomes even clearer when we consider those (admittedly rare) cases where it is not by acting on impulses, but by *suppressing* them that coordination is achieved. As any continental European or American who has ever made the experience of driving a car in Great Britain or Australia will confirm, it is possible to coordinate along the lines of the left hand traffic rule in spite of persisting impulses to the contrary. In the case of the people coming from continental Europe driving off the car ferry in Dover, a great deal of suppression of impulses is required. It seems plausible that this is only possible because under normal circumstances (i.e. where common knowledge of those rules and the absence of suicidal preference can be assumed), people such as student drivers and foreigners find it perfectly *rational* to stick to the traffic regulations in order to avoid collisions.

To this an irrationalist philosopher of coordination might object the following: the fact that under special and unusual circumstances such as when learning to drive, or when driving off the car ferry in Dover, coordination is achieved rationally without appropriate impulses, or even by suppressing one’s impulses, does not prove against the fundamental role of impulses in coordination, because it is *the other drivers’ impulses*, to which it is rational to adapt one’s

own behavior. Thus it still seems to hold that any sound reasoning about what equilibrium to aim at in coordination games ultimately bottoms out in mere impulses. However, this defense of the irrationalist position does not stand up to closer scrutiny. It is not necessary either to act on impulses or to count on the other's appropriate impulses to achieve coordination.

Coordination can even be achieved where both parties act against their impulses, and where this is common knowledge.

The following example may serve as an illustration. The island of Jersey, the largest of the Channel Islands, is popular with tourists, many of which come over from the continent, since Jersey is only at some fourteen miles off the French coast. There is public transportation on the island, but in order to avoid the crowd and to reach the most beautiful places, many tourists either rent a car on the island, or have their own car brought over by ferry, sometimes even just for a day-trip. There are many narrow roads with no separated lanes on Jersey; in order to avoid the branches sticking out of the hedges on both sides, most cars drive in the middle of the road, moving aside (often without reducing their considerable speed) only to let oncoming cars pass by. Because the states of Jersey are a part of the United Kingdom, this is done by both swerving to their left. Drivers who come from the continent have, of course, to suppress their impulses in order to adapt to the left-side traffic rule. What flies directly in the face of the irrationalist position, however, is this: even drivers that are clearly recognizable to each other as coming from the continent by their number-plates and by the location of the driver's seat in their cars coordinate without any difficulty by keeping to their left.

From an irrationalist viewpoint, this must appear like a small miracle. For clearly, these drivers neither act on impulse, nor do they rely on the impulses of the other drivers (since it is common knowledge between them that both drivers are from the continent). But if it is true that in this particular situation, there are indeed no appropriate impulses, habits, psychological propensities and such around: how then is coordination possible? The obvious answer is very simple: because contrary to what the irrationalists say, common sense is right: Given common knowledge of the rules and of the absence of suicidal preferences, sticking to the basic traffic rules is simply the *rational* thing to do.

It is true that as experienced drivers in our everyday world, we do not *think* about whether or not we should stick to the rules; we just do it "blindly". This however does not mean that rationality is not involved here, or that it comes second to our habits and impulses. For if we let ourselves "blindly" be guided by the rules, we do this precisely because we think that this is the rational thing to do. And how could this belief be so pervasive if it was all wrong?

#### IV. Team Thinking

If the foregoing is right, i.e. if we can neither accommodate our pre-theoretic intuitions in our standard theory of rationality in action (II.), nor let go of our pre-theoretic convictions concerning the rationality of coordination (III.), it seems that the correct position will be the only one that is left: to do something about the theory in order to make it fit our deeply engrained pre-theoretic conviction. If so little can be said against our driver from within the conceptual framework of our standard theoretic understanding of rationality in action, we will have to revise this framework. Obviously, there is something more to rationality in coordination than mere individual expected utility maximization in the sense discussed above.

There are several theories that point the way to go. In their *General Theory of Equilibrium Selection in Games*, John Harsanyi and Robert Selten claim that the principle of coordination – their term is “payoff dominance” – *cannot be derived from individual rationality*, but implies an altogether different, a *collective* concept of rationality:

„Our theory uses two independent, and ostensibly very different, criteria of rationality. One of them, risk dominance, is based on *individual* rationality: it is an extension of Bayesian rationality from one-person decisions to n-person games involving strategic interaction among n players, each of them guided by Bayesian rationality (...). In contrast, payoff dominance is based on *collective* rationality: it is based on the assumption that in the absence of special reasons to the contrary, rational players will choose an equilibrium point yielding all of them higher payoffs, rather than one yielding them lower payoffs. That is to say, it is based on the assumption that rational individuals will cooperate in pursuing their common interest if the conditions permit them to do so“ (Harsanyi/Selten 1988: 365).<sup>6</sup>

Picking up on Harsanyi’s and Selten’s insight, Robert Sugden has developed his theory of *team thinking* (Sugden 1993; 1996; 2000; 2003). Very roughly, the basic idea of Sugden’s theory as well as of Michael Bacharach’s somewhat related account (Bacharach 1998) seems to be the following. The standard theory’s fault is that it conceptually restricts the “units of agency” to single individuals. This leads to an inadequate account of those situations, where we do not reason and act as single isolated individuals, but as *members of teams* instead. “Team membership” is basically meant in the sense of participation in collective action. To

---

<sup>6</sup> Raimo Tuomela seems to adopt a similar view on the deconditionalization problem. He says that deconditionalization is not a “fully ‘r-rational’” procedure (Tuomela 2002: 395f.), where “r-rational” is something like individual instrumental rationality. However, he distinguishes r-rationality from a wider everyday sense of rationality, which seems to include the possibility of rational deconditionalization.

understand the structure of *team thinking*, it is important to see the situation at hand as one of *shared intentionality*.

Here, the recent turn to collective intentionality comes into play. Whereas ‘classical’ philosophy of mind focused exclusively on the analysis of *individual* intentionality, it has become increasingly clear over the last two decades that in order to account for the *social dimension of human action and cognition* the analysis has to be extended to *shared* intentional states. Based on seminal contributions dating from the eighties, the analysis of *collective intentionality* has gradually evolved into a distinct field of research in Europe and America. The most important theories of collective intentionality are those by Raimo Tuomela (1988; 1995), Margaret Gilbert (1992), John Searle (1995), and Michael Bratman (1999). These theories differ in very fundamental ways. Tuomela’s account rests on individuals’ intentions to do their part together with a structure of mutual belief. Searle criticizes Tuomela and claims that collective intentions are irreducible to sets of individual intentions. Searle’s account rests on intentions of the form “we intend...” in the individual minds of the participants. Bratman, on his part, gives an account of collective intentions in terms of interrelations of individual intentions of the form “I intend that we...”. Gilbert again follows an altogether different line by making *collective commitments* the center of her account of collective intentions. These differences and open issues<sup>7</sup> notwithstanding, the importance of collective intentionality analysis for our understanding of both the mind and the social world has been widely recognized in philosophy as well as in many neighboring disciplines.

Can collective intentionality analysis indeed help to understand the rationality of coordination? With regard to the example of our drivers, this might seem rather unlikely at first. For obvious reasons, the paradigm cases of shared intention are cooperatively loaded cases such as John Searle’s example of the joint intention to cook a sauce hollandaise (Searle 1990: 400ff.), Margaret Gilbert’s example of the joint intention to go for a walk (Gilbert

---

<sup>7</sup> The main controversies in the current debate on collective intentionality revolve largely around the following three key issues: 1. *The Intentionality Issue*: What is the structure of the *intentionality* involved in collective intentional states? Is the difference between collective intentionality and individual intentionality primarily a matter of the *content*, the intentional *mode* or rather a matter of the *subject* of the mental states in question? 2. *The Coordination and Commitment Issue*: What is the “binding force” of collective intentional states? How, if at all, do *commitments* and *obligations* come into play? Does collective intentionality analysis require us to depart from the existing approaches to human behavior? 3. *The Collectivity Issue*: In what sense is *collectivity* involved in collective intentionality? Does collective intentionality presuppose (or constitute) collectives or supra-individual agents/subjects?

1996b) or Raimo Tuomela's example of a group of people joining their forces to push a broken-down car (Tuomela 1995: 137f.). By comparison to such intensely communal endeavors, it might seem that there is *nothing genuinely collective or social* our drivers' intention to avoid bumping into another driver whom he does not even know by sticking to the right side traffic rule.

The theory of *team thinking*, however, points out the hidden element of sharedness that is implicit in these cases. If "avoiding a collision" is seen as something the two drivers desire *individually*, the deconditionalizing problem appears to be insoluble. If we conceive of the participating individuals as acting on their *individual* desire to get by the other collision free, we immediately find ourselves caught up in the regress of interdependent expectations that our driver pointed out. This is not the case, however, for drivers who are seen as basically *sharing* the aim to get by each other collision-free. Given the fact that what we *together* intend to do is to get by each other, what is rational for *me* is to perform my share of what maximizes our shared desire. Thus "right" is not rational because it immediately yields a better outcome (which it does only *conditionally*; if the other driver were to choose "left", my choosing "left" would be in the best interest not only of myself, but of the team, too). It is rational for me to choose "right" because it is my part in what *we* should be doing.<sup>8</sup> The fact that individuals can be team members has consequences for what is rational for them to do. For "one of us", the decision to move left is plainly irrational. Thus the deconditionalizing problem is solved.

The theory of team thinking requires a theory of rationality, intentionality, and action, that is richer as the one that is implicitly adopted in the standard economic model of behavior, because it allows for a sense in which *teams* can be said to have preferences (or even make choices) which are, in a certain sense, *irreducible* to simple individual preferences (or choices) (cf. Sugden 2000). In this view, not all preferences, goals, and desires, and other intentional states, are individual goals and intentional states. In the case of the driver's coordination problem, there are not two separated individual goals not to collide, but the

---

<sup>8</sup> In Susan Hurley's (1989: 145ff.) view, "right" is not rational because of its *causal* consequences, but because of its *constitutive* consequences, i.e. because our individual choosing "right" *constitutes* the action that is best for both of us. I am somewhat uncomfortable with Hurley's way of rationalizing coordination, because it seems obvious that it renders the question "should I do what constitutes my part in what constitutes the best collective choice, or should I rather do my part in what constitutes the second best collective choice" unanswerable. Because obviously, it is rational to choose the latter alternative *if* one expects the other to choose that alternative, too. In other words: Hurley's *constitutive rationality* does not solve the deconditionalizing problem.

participants act on a *shared* goal instead (cf. Schmid 2005a). Getting by each other collision-free is not anything the single individual drivers want. It is something they want *together*. This desire is *irreducible* in the sense that it is not the case that the drivers share their desire (have a preference for “right”/”right”) because they have the appropriate individual desires (i.e. an individual preference for “right”); rather, their individual “contributive” intentions or preferences are *derived* from the shared intention or preference.

Thus in this situation, the drivers do not appear as distinct units of agency, but as members of a whole that in a sense appears to be capable of thinking and acting. To capture this trait of team thinking, Sugden invokes a rather strong concept of collectivity. It seems that there are not only individuals at the basic level of explanation of social phenomena, but also teams, to which, following Sugden, the participating individuals are members “in something like the old sense in which arms and legs are members of the body” (Sugden 1993: 86).

The idea of some irreducible sense of collectivity goes much against the individualistic grain of current social theory and social science. It might even appear that in the theory of team thinking, some somber group mind raised its head. Indeed it seems that Sugden himself loses some of his anti-individualistic courage when getting sight of these possible ontological consequences of his theory. What could possibly save us from ending up in a collectivist group mind conception once we start loosening the individualistic restriction of the classical account? Sugden resorts to the following solution. In a rather harsh contrast to his strong concept of membership, he hastens to assert that the existence of the collective depends on the *participating individual’s thinking of themselves in terms of team members* (which conforms to the classical individualistic Weber-style view of the social, in which collectivities are “real” only as parts of the contents of the intentional states of individuals).<sup>9</sup>

In a much similar vein, Michael Bacharach based his theory of *Team Reasoning* in some “group identification”, in which individuals come to take themselves to be members of a team: „in certain circumstances, individuals tend to identify themselves with a group; and a group identification leads them to team-reason“ (Bacharach 1998: 132). Besides these two important philosophers, there are other attempts to reconcile the acknowledgement of the role of some kind of “team reasoning” with an individualistic ontology of action. Thus Maarten C.W. Janssen puts forth an even weaker version of the role of collectivity in coordination, replacing “collective rationality” with what he calls “individual team member rationality”. Again, the ontological line behind this is stoutly individualistic: „Where there is enough infor-

---

<sup>9</sup> „A team exists to the extent that its members take themselves to be members of it“ (Sugden 2000: 192).

mation and knowledge about each other, players can consider themselves as a team and think individually what is best for the team and its members“ (Janssen 2000: 13).

There are at least two reasons, however, to reject the view that team reasoning (collective rationality, team thinking, individual team member rationality or whatever it might be called) depends on some “taking oneself to be a member of the team” from the side of the individual members (or some reflective “group identification”, or some such). First, this view seems rather absurd if we consider cases such as the one of our drivers. We obviously do not have to “take ourselves to be members of a team” to find it rational to stick to the traffic regulations in order to avoid accidents. If team thinking is at work in these cases (and I believe that it is), the element of collectivity involved here is obviously *not a matter of some reflective attitude*, or belief about oneself, for it seems that phenomenologically speaking, there are no such attitudes whatsoever involved here. Team thinking does not require that the participating individuals take themselves to be members of the team.<sup>10</sup> Reflective awareness of one’s status as a team member is neither a sufficient nor a necessary condition of team thinking. It is not sufficient, because one can mistakenly identify oneself as a team-member. And it is not necessary, because one can be a member of a team without reflectively identifying oneself as such. In this sense, team thinking is *pre-reflective*. This also means that if we do correctly think of ourselves as members of some team, this is because *we are* a team, and not the other way around, as Sugden and those subscribing to a similar view of the role of collectivity in coordination seem to believe.

The second reason is the following. Consider again the driver’s coordination problem. From the viewpoint of the theory of team thinking, team reasoning, individual team member rationality, or some such, whether or not “right” is rational ultimately depends on *who one takes oneself to be*. If I take myself to be “one of us” (“the other driver and me”), i.e. a member of a team, “right” is the rational choice, because “right”/”right” is what *we* intend to achieve, and my choosing “right” is my individual contribution to our shared goal, whereas for somebody who exclusively optimizes her or his individual desires, “right” is just as hypothetically rational as “left”. It all depends on one’s identity in terms of one’s reflective understanding of oneself. Identity, one could say, is a matter of self-categorization, and it is prior to rationality.

---

<sup>10</sup> This is obscured in Sugden’s account because of Sugden’s preoccupation with cooperatively more loaded cases such as his Footballer’s coordination problem (Sugden 2000: 2003).

Thus these theories of *team thinking*, *team reasoning*, *individual team member rationality* and so on seem to offer a kind of a compromise between our driver's way of reasoning on the one hand and common sense on the other in that they hold on to the commonsensical view that it is rational to choose "right", while at the same time, some tacit assumption that is accepted throughout the debate is jettisoned. Why should it be necessary to be able to prove our driver wrong, if one claims that "right" would have been the rational choice? Sugden, Bacharach and Janssen seem to *maintain* the commonsensical claim to rationality without having to bear the burden of proof against our driver. "Right" is the rational choice – from the perspective of a team member, not from the perspective of an isolated individual. What is rational to do depends on who one takes oneself to be.<sup>11</sup>

As convincing this relativizing move might seem with respect to the trouble with the *rational fool* of our initial example, there are some serious doubts left. Let us again take the pre-theoretical, commonsensical perspective. If in our everyday-understanding, we take "right" to be the rational choice, we take "right" to be the rational choice, full stop. We do not mean something like "'right' is the rational choice *for people who take themselves to form a spontaneous team together with the oncoming drivers*", or "'right' is the rational choice for people who do not only individually prefer to avoid a collision, but *team-prefer* to get by the other". If we call our driver *irrational* given all he says about the circumstances of his decision, we do not mean something like "irrational as a team-member, but rational from the perspective of an atomistic individual". Once again, we mean *irrational*, full stop. Thus it seems that if indeed team thinking is involved in coordinative behavior of this kind, and if we are right in holding on to the commonsensical view concerning the rationality of such behavior, a more robust conception of the collectivity of team thinking than the one put forth by Sugden is needed. No matter what our self-image might be: we simply *are* team members in these situations, and as such we share our intentions and goals. What makes our driver irrational is that he is not aware of who he is: not an isolated individual, but *one of us*.<sup>12</sup> The individualistic social ontology to which these theorists subscribe has to be dropped.

---

<sup>11</sup> For a clear statement of this cf. Elizabeth Anderson's "Priority of Identity to Rational Principle": „what principle of choice it is rational to act on depends on a prior determination of personal identity, of who one is. The validity of the principle of expected utility (maximizing the satisfaction of one's personal preferences) is *conditional* on regarding oneself as an isolated individual, not a member of any collective agency“ (Anderson 2001: 30).

<sup>12</sup> As argued above, this does not mean that we have to be *reflectively aware* of our "true" nature as social beings in order to avoid collisions in everyday life. Indeed the fact that even individuals who take themselves to be

This, however, is not easily done. For as we shall see in the next section, it seems that dropping individualism about intentionality is tantamount to endorsing collectivism: the group mind raises its ugly head. Fear of the group mind is so pervasive in the current debate, and plays such an important role in the individualistic setting that most philosophers of collective intentionality endorse, that I think it is justified to spend the last section of this paper trying to address these worries. This is not only to postpone the burdensome task of delineating a more adequate (i.e. less individualistic) account to future explorations of the issue. Rather, it is in the light of these considerations that two further basic features of a more robust conception of collective intentionality are revealed. Collective intentionality is not just pre-reflective. It is also, as will be argued in the following section, both *irreducible* and *relational*.

#### V. Who's Afraid of the Group Mind?

What seems unfortunate in the theory of *team thinking*, as well as in similar accounts, is that the existence of the team is made a matter of its member's taking themselves to be its members. However, it was only with the seemingly honorable aim to avoid ending up in some collectivist group mind conception that the theory of team thinking resorted to basing the existence of the team in some reflective individual attitude in the first place. Thus the decisive question is: does a stronger account of collective intentionality imply endorsing some collectivist macro-subject, or group-mind?

This is a question that is of great importance to the current debate on collective intentionality in its entirety. For here, the specter of the group mind is ubiquitous. Its prominent role in this debate seems to stem from a rather innocent-looking assumption. Where there is intentionality, it is said, there has to be somebody who 'has' it – the good old subject. Now if it is claimed that there is such a thing as *collective* intentionality, and that collective intentionality has to be distinguished from *individual* intentionality, the conclusion seems to force itself on us that it has to be not the single individuals, but the collectives themselves that 'have it'. And for collectives to have intentions, some sort of a 'collective mind', some 'group mind' seems to be required, something hovering over and above the mind of the individuals involved.

This apparent implication of the very concept of collective intentionality does not look very appealing. Even though there is no consensus concerning the structure of collective

---

atomistic *homines oeconomici* can successfully coordinate in real life shows how far our reflective self-image can depart from our pre-reflective way of reasoning.

intentionality, all philosophers of collective intentionality agree that there is no mind over and above the minds of individuals. Might the question of whether or not (and if so, in what sense) collectives can act remain to some degree controversial, it seems obvious that it is unacceptable to treat collectives as ‘subjects’ of intentions and actions in the ordinary sense in which individuals are the bearers of their intentional states. Even if the notion of the collective subject is stripped of its mentalistic content, it still does not quite appeal to us because it is vaguely associated with collectivistic (or even totalitarian) notions of the social. If the collective ‘has’ intentionality, the individual seems to be no more than an organ, i.e. merely an instrument, which flies in the face of our idea of individual intentional autarky or autonomy (cf. Pettit 1996: 117ff.).

Thus it seems quite understandable that all philosophers of collective intentionality set themselves the task to show that collective intentionality is possible without there being a group mind involved. The specter of the group mind (or collective subject) has to be exorcised, and one can identify two different ways in which this is done. The softer way is the one chosen by Gilbert (1989), and Tuomela (1995: 231). Some sort of collective subject is admitted to the theory (in Gilbert’s case the *Plural Subject*), but it is domesticated as to be consistent with an otherwise thoroughly individualistic conceptual framework. The collective subject is solidly founded in the intentional autonomy of individuals by reducing the collective subject either to sets of individual intentions (Tuomela 1995) or to the reflective self-understanding of the single participating individuals as members of the team (cf. Gilbert 1992: chap. 4). The tougher way of dealing with the specter of the group mind is simply to treat it as an absurd collectivist idea that has to be banished from the theory of collective intentionality straight away – Searle sometimes speaks of „a perfectly dreadful metaphysical excrescence“ (Searle 1998: 150; cf. also Searle 1990: 406; for Bratman’s view on the collective subject see Bratman 1999: 111; 122f.).

On this ‘hard’ line against the collective subject, the group mind is exorcised in two different ways involving two different kinds of individualism – subjective and formal individualism –, one of which is rejected while the other is endorsed. Searle makes a bold move beyond *formal* individualism by claiming that collective intentionality is *irreducible* to individual intentionality. Collective intentionality consists of intentions of the *form* “We intend...”. In order to exorcise the group mind, however, he endorses *subjective* individualism by emphasizing that collective intentionality is exclusively *in the minds of individuals* and independent of anything external (Searle 1990: 406f.).

Interestingly, Bratman seems to make the opposite move. Going beyond *subjective* individualism, Bratman emphasizes the *relational* character of shared intentionality. As “interlocking webs” (Bratman 1999: 9) of the participant individuals’ intentions, shared intentionality is structurally dependent of entities external to the individual mind of each single participant. At the same time, however, and in a certain tension with this bold interrelationalistic move beyond subjective individualism Bratman exorcises the group mind by stressing that his account is thoroughly “reductive in spirit” (ibid.: 108) and endorsing formal individualism.<sup>13</sup>

Looking from afar at how the group mind is dealt with it might appear that the theory of collective intentionality is caught in a dilemma, or rather, stuck in some kind of Double Bind. On the one hand, the aim is to break with individualism in the sense of the orthodox concentration on purely individual intentionality. On the other hand, individualism (in the broad sense of an emphasis on the role of the individual) seems to be the only effective means against the group mind. And this is what causes the most difficulties in these two influential accounts of collective intentionality. Searle has been rightly criticized for his subjective individualism, while the focus of critique in the debate that followed Bratman’s account is on his formal individualism (cf. Schmid 2003b for a detailed discussion).

I do not have the space to present the arguments here, but the upshot of this is that an adequate account of collective intentionality should not depart *either* from subjective *or* from formal individualism. Rather, it should depart from *both* versions of individualism. Instead of arguing directly for a non-individualistic account of collective intentionality, let me turn to the most obvious objection any such account will have to face instead. If both formal and subjective individualism is left behind, what, then, about the group mind? How can a robust conception of collective intentionality avoid flying into the face of our intuitive notion of individual intentional autarchy?

I believe that the answer is rather simple. The whole trouble with the group mind arises from the attempt to assign collective intentionality to *one* individual source or bearer. It arises from the attempt to give some acceptable answer to the question: who is it that *has* collective intentions? Who is it that *does* the team thinking? And this question, innocent as it might look, is heavily loaded with historical ballast that we should, simply jettison and leave behind. Only in the last decades, we have successfully managed to get rid of Descartes’ quest for absolute certainty in philosophy. However, the Cartesian preoccupation with the individual

---

<sup>13</sup> Cf. Bratman’s conceptual analysis in Bratman (1999), p. 105.

subject still persists. It is still a deeply rooted idea that where there is intentionality, there has to be *one* somebody who “has” it as its owner, source or bearer – if not the individual human beings, than some individual group mind. It is the fact that most philosophers of collective intentionality hold on to this assumption that gives rise to the fear that by moving too far away from individualism, we are running the risk of getting stuck with the group mind. Collective intentions, however, are not intentions of the kind anybody *has* for herself – not single individuals, and not some group mind. Rather, it is something individuals *share*. Indeed, there is all the less reason to resort to individualism in order to exorcise the group mind as the latter is clearly itself an individualistic concept. The group mind arises from the conceptual compulsion to assign a single subject to collective intentions. Collective intentions, however, do not have a *single* subject. They have many. Thus the group mind is nothing we should be afraid of. It is merely a distorted individualistic image of a non-individualistic, holistic concept of the mind.

## *VI. Conclusion*

In this paper, I have argued that collective intentionality permeates human interaction down to its very basic modes. It is our capacity for collective intentionality that deconditionalizes rational decision-making, where coordination among agents is required. Where there is no incentive for unilateral defection, rational agents will think and act as a team. In order to account for this, however, it is necessary to depart from individualistic assumptions such as implied in the standard economic standard model of behavior, and to widen the perspective to collective intentionality.

Secondly, I have argued that a more robust conception than the one to be found in the received literature is necessary in order to develop an adequate account of collective intentionality. I have introduced the following three features of a robust conception (though only the first feature is developed in some detail above):

- 1) Collective intentionality is *pre-reflective*. It is not a matter of some reflective attitude of the participating individual (i.e. the individual’s taking herself to be a member of a team), or “self-categorization”.
- 2) Collective intentionality is *irreducible* to individual intentionality, i.e. it does not consist in some set of intentions of the form “I intend...”. An adequate account is incompatible with *formal* individualism.

- 3) Collective intentionality is *relational*. An adequate account is incompatible with *subjective* individualism.

In the last section of the paper, I have claimed that widespread fears that a more robust conception of collective intentionality will get stuck with the group mind are mistaken. What are the consequences of collective intentionality analysis for economic theory? The widening of perspective as implied in collective intentionality analysis directly affects the notion of the agent. This meets with other tendencies in “heterodox” economic theory. As John B. Davis has pointed out in his book on the theory of the individual in economics (Davis 2003: esp. pp. 130–149), collective intentionality analysis seems to mesh seamlessly with an increasing unease with the “atomistic” standard model of the agent. In this vein, collective intentionality analysis is particularly attractive because it opens a perspective on social identity and human embeddedness that does not hinge on adventitious stigmata such as birth and destiny. There is a tendency in the received literature to conceive of social identities as fixed entities. As Amartya K. Sen (2004) has convincingly argued, however, social identities are *made* rather than *discovered*, they are a matter of what we *do* rather than a matter of *what we are*. For all of the work that needs to be done in this relatively new field of research, collective intentionality analysis seems to be a promising candidate to be able to show how these identities come about.<sup>14</sup> If the argument developed in this paper is right, it seems that some very rudimentary forms of “social identity” – i.e. shared goals, pursued by a team – are at play even in the most transient of our interactions, such as the one of two drivers successfully passing by each other on the highway.

### *References*

- Anderson, Elizabeth (2001) ‘Unstrapping the Straitjacket of ,Preference’: A Comment on Amartya Sen’s Contributions to Philosophy and Economics’, *Economics and Philosophy*, 17: 21-38.
- Bacharach, Michael (1998) ‘Interactive Team-Reasoning: A Contribution to the Theory of Co-operation’, *Research in Economics*, 58: 117-147.

---

<sup>14</sup> For an interpretation of Sen’s influential criticism of orthodox economic theory in terms of collective intentionality analysis see Anderson 2001; Schmid 2005a. For a discussion of the importance of collective intentionality analysis for experimental economics cf. Schmid 2005b.

- Bratman, Michael E. (1999) *Faces of Intention. Selected Essays on Intention and Agency*, Cambridge: Cambridge University Press.
- Davis, John B. (2003) *The Theory of the Individual in Economics*, London: Routledge.
- Gauthier, David (1975) 'Coordination', *Dialogue*, 14: 195-221.
- Gilbert, Margaret (1992) *On Social Facts*, Princeton: Princeton University Press.
- Gilbert, Margaret (1996) 'Rationality and Salience', in *Living Together. Rationality, Sociality, and Obligation*, Lanham: Rowman and Littlefield, pp. 23-37.
- Gilbert, Margaret (1996b) 'Walking Together', in: *Living together. Rationality, Sociality, and Obligation*, Lanham: Rowman and Littlefield, pp. 177-194.
- Goyal, Saanjev and Janssen, Maarten Chr. (1996) 'Can we rationally learn to coordinate?' *Theory and Decision*, 40: 29-40.
- Harsanyi, John C. and Selten, Robert (1988) *A General Theory of Equilibrium Selection in Games*, Cambridge Mass.: MIT Press.
- Hurley, Susan (1989) *Natural Reasons. Personality and Polity*, Oxford: Oxford University Press.
- Janssen, Maarten C. W. (2000) 'Towards a Justification of the Principle of Coordination', *Tinbergen Institute Discussion Papers* Nr. 00-017/1.
- Janssen, Maarten C. W. (2001) 'On the Principle of Coordination', *Economics and Philosophy*, 17: 221-234.
- Janssen, Maarten C. W. (2001b) 'Rationalizing Focal Points', *Theory and Decision*, 50: 119-148.
- Lewis, David K. (1969) *Convention. A Philosophical Study*, Oxford: Blackwell.
- Miller, Seumas (1991) 'Co-ordination, Salience, and Rationality', *The Southern Journal of Philosophy*, 29/3: 359-371.
- Parsons, Talcott and Shils, Edward A. (1959) 'Categories of the Orientation and Organization of Action' in T. Parsons and E. Shields (eds) *Toward a General Theory of Action*. Cambridge Mass., pp. 53-109.
- Pettit, Philip (1996) *The Common Mind. An Essay on Psychology, Society, and Politics*, New York: Oxford University Press.
- Provis, C. (1977) 'Gauthier on Coordination', *Dialogue*, 16: 507-509.

- Schelling, Thomas C. (1960) *The Strategy of Conflict*, Cambridge Mass.: Harvard University Press.
- Schmid, Hans Bernhard (2003a) 'Rationality-in-Relations', *American Journal of Economics and Sociology*, 62/1, Special Invited Issue on John R. Searle's Philosophy, 67-101; reprinted in Laurence S. Moss and David Koepsell (eds) (2003) *John Searle's Ideas About Social Reality. Extensions, Criticisms, and Reconstructions*, Oxford: Basil Blackwell).
- Schmid, Hans Bernhard (2003b) 'Can Brains in Vats Think as a Team?' *Philosophical Explorations*, 6/3: 201-218.
- Schmid, Hans Bernhard (2005a) 'Beyond Self-Goal Choice. Amartya Sen's Analysis of the Structure of Commitment and the Role of Shared Desires', *Economics and Philosophy*, 21/1: 51-63.
- Schmid, Hans Bernhard (2005b) 'Nostrism – Social Identities in Experimental Games', *Analyse & Kritik*, 27/1, Special Issue: Ernst Fehr on Human Altruism, 172-187.
- Searle, John R. (1990) 'Collective Intentions and Actions', in Philip R. Cohen, Jerry Morgan and Martha E. Pollack (eds) *Intentions in Communication*, Cambridge Mass.: MIT Press, pp. 401-415.
- Searle, John R. (1995) *The Construction of Social Reality*, New York: Free Press.
- Searle, John R. (1998) 'Social Ontology and the Philosophy of Society', *Analyse & Kritik*, 20: 143-158.
- Segal, Gabriel M.A. (2000) *A Slim Book about Narrow Content*, Cambridge Mass., MIT Press.
- Sen, Amartya K. (2004) 'Social Identity', *Revue de philosophie économique*, 7: 7-27.
- Sugden, Robert (1993) 'Thinking as a Team: Towards an Explanation of Nonselfish Behavior', *Social Philosophy and Policy*, 10: 69-89.
- Sugden, Robert (1995) 'A Theory of Focal Points', *The Economic Journal*, 105: 533-550.
- Sugden, Robert (1996) 'Rational Coordination', in F. Farina, F. Hahn and S. Vanucci (eds) *Ethics, Rationality and Economic Behavior*, Oxford: Clarendon Press, pp. 244-262.
- Sugden, Robert (2000) 'Team Preferences', *Economics and Philosophy*, 16: 175-204.

Sugden, Robert (2003) 'The Logic of Team Reasoning', *Philosophical Explorations*, 6: 165-181.

Hollis, Martin and Sugden, Robert (1993) 'Rationality in Action', *Mind*, 102: 1-35.

Thalos, Mariam (1999) 'Degrees of Freedom: Towards a Systems Analysis of Decision', *Journal of Political Philosophy*, 7: 453-477.

Tuomela, Raimo and Miller, Kaarlo (1988) 'We-Intentions', *Philosophical Studies*, 53: 367-389.

Tuomela, Raimo (1995) *The Importance of Us. A Study of Basic Social Notions*, Stanford: Stanford University Press.

Tuomela, Raimo (2002) 'Joint Intention and Commitment', in Georg Meggle (ed.) *Social Facts and Collective Intentionality*, Frankfurt am Main: Hänsel-Hohenhausen, pp. 385-418.