

Epistemic Conditions for Collective Action

SARA RACHEL CHANT AND ZACHARY ERNST

Writers on collective action are in broad agreement that in order for a group of agents to form a collective intention, the members of that group must have beliefs about the beliefs of the other members. But in spite of the fact that this so-called 'interactive knowledge' is central to virtually every account of collective intention, writers on this subject have not offered a detailed account of the nature of interactive knowledge. In this paper, we argue that such an account is necessary for any adequate analysis of collective intention. Furthermore, we argue that an application of Robert Aumann's theory of interactive knowledge may be used to address several puzzling features of collective intention.

Introduction

Actions undertaken by groups often require the epistemic states of the agents in the group to depend appropriately on the epistemic states of other agents in the group. For example, we may successfully coordinate our meeting only on the condition that we both know where the meeting is to take place, and that we know the other person knows this. Such situations are not unusual—there are important classes of cases in which interdependence among our beliefs is precisely what marks the difference between a collective action and an arbitrary set of actions performed by individuals.

Thus, Michael Bratman (1992, 1993), Margaret Gilbert (1990), Christopher Kutz (2000), and Raimo Tuomela (1984, 1991), have each suggested that in order for a group to perform a collective action, the members of that group must satisfy certain epistemic conditions.¹ Tuomela, for example, argues that acting together may require higher-order 'loop beliefs' in which I believe that you will participate in the performance of the act, I believe that you believe that I will participate, and so on (Tuomela 2002, p. 28). If such higher-order beliefs are indeed

¹A variety of different terms referring to such actions have been adopted in the action theory literature, including 'collective action', 'joint action', 'social action', and 'shared cooperative activity'. Although these different terms often have different meanings, we shall refer to any action brought about by a group of agents, and requiring coordination among those agents, as a 'collective action'.

necessary for the performance of collective actions, then it is important to understand the conditions that make those epistemic states possible. In this paper, we shall concentrate on so-called ‘interactive knowledge’, which refers to the knowledge possessed by one agent about the knowledge possessed by another agent.

In work that deserves more attention from philosophers, the economist Robert Aumann has developed a formal theory of interactive knowledge (Aumann 1999a, 1999b). Although some aspects of his theory are controversial, we shall argue that it provides a useful framework for thinking about interactive knowledge and its role in producing collective action. In section 1, we briefly review the most influential accounts of collective intention, highlighting the central role played by interactive knowledge in these accounts. Section 2 considers a case in which rational agents can never coordinate their actions despite achieving arbitrarily high degrees of interactive knowledge. This case, known as the ‘Coordinated Attack Problem’, will motivate our argument that there is no simple, common-sense criteria for specifying what level of interactive knowledge is necessary for collective action. In section 3, we briefly sketch Aumann’s theory of interactive knowledge. In section 4, we argue that the level of interactive knowledge required for collective action depends on the strategic situation in which the agents are supposed to act. In fact, we shall see that there are counterintuitive (but realistic) situations in which interactive knowledge *prevents* collective action, rather than enabling it. Section 5 discusses an objection that has been made against Michael Bratman’s theory of ‘shared cooperative activity’. We argue that Bratman’s response, while sound, can be strengthened and generalized if we take into account how interactive knowledge is generated by a group. The final section of the paper briefly outlines a few open problems that would benefit from a similar approach.

1. The importance of interactive knowledge in philosophical action theory

A common strategy for analysing the concept of collective action is to assume that just as individual intentions play an important role in the explanation of individual action, so too collective intentions play an important role in the explanation of collective action. But the concept of a collective intention is far from clear. Accordingly, there is little agreement among the analyses as to the exact nature of collective intentions or whether collective intentions exist in any meaningful sense at all. However, there is an important point upon which most accounts of

collective intention agree—namely, that collective intentions require that the agents' beliefs depend in some appropriate way on the other agents' beliefs.²

Michael Bratman's analysis of collective intentions is representative of this approach. In summarizing his view of collective intentions, Bratman writes:

[Collective intentions] are interlocking in the sense that we each intend both that our *J*-ing proceed by way of each person's relevant intentions and that these intentions appropriately mesh. I also supposed that this interlocking web of intentions would be out in the open, a matter of common knowledge.³ (Bratman 1999, p. 143)

This requirement that the 'web of intentions' be 'out in the open' is something that Bratman has elaborated upon in various places. In his article, 'Shared Cooperative Activity', he says that:

In [shared cooperative activity] the fact that there is this mutually uncoerced system of intentions will be in the public domain. It will be a matter of common knowledge among the participants. I will know that we have these intentions, you will know that we have these intentions, I will at least be in a position to know that you know this, and so on. (Bratman 1993, p. 335)

Other writers on collective action have introduced similar requirements on collective intention. To take another example, Christopher Kutz has introduced a requirement that he calls 'mutual openness'. This requirement is motivated in much the same way as is Bratman's common knowledge requirement. In discussing an example in which two agents wish to coordinate their plans, Kutz says that:

... for our going to Chicago together to be joint, we each must believe it at least possible the other knows of or will try to predict our choice, and be favorably disposed to the other's knowledge of anticipation of that choice ... our individual aims are furthered rather than frustrated by the other's awareness. (Kutz 2000, p. 6)

On Kutz's view, the requirement of 'mutual openness' is supposed to be weaker than Bratman's common knowledge requirement, because there are cases in which it may be possible for two agents to coordinate their

² The most important exception to this generalization is the theory advocated by John Searle, who argues that collective intentions are primitives in the heads of individual agents (Searle 1990 and 1997). Indeed, it is an implication of his theory that a person might have a collective intention even though there are no other individuals with whom she can share that intention. For the purpose of this paper, however, we shall discuss only the more common strategy of assuming that there must be some appropriate relationship between the beliefs of the various agents in the group.

³ It is important to note that Bratman explicitly says that he is not using the term 'common knowledge' in the precise sense first defined by David Lewis (1969).

actions without knowledge of what the other agent is going to do—as in cases where the right choice is highly salient to both parties. However, in spite of the weakness of this requirement, it is still the case that some sort of interactive knowledge is taken to contribute to the formation of a collective intention. Interactive knowledge ‘furthers’ rather than ‘frustrates’ the agents’ attempts to coordinate their actions.

Raimo Tuomela’s analysis of collective intentions also contains a similar belief requirement. According to Tuomela, a group of agents has a collective intention if and only if the members of the collective have the right kind of individual intention, which he has termed a ‘we-intention’. And in order for an individual to have a we-intention, she must ‘believe that there is ... a mutual belief among the participating members of [the group] ... to the effect that the joint action opportunities for an intentional performance of X will obtain ...’ (Tuomela 1991, p. 265). Here, Tuomela uses the term ‘mutual belief’ in the sense that there is a mutual belief that p if everyone in the group believes that p . Thus, Tuomela’s account demands a certain kind of interactive knowledge in requiring that an agent believe that there is a mutual belief among the members of the group.

2. The Coordinated Attack Problem

One might suspect that the details of interactive knowledge do not matter to an analysis of collective intentions. For example, one might suppose that there is some requisite threshold level of interactive knowledge—perhaps vague—at which collective intention (and thereby, collective action) becomes possible. If this were true, then writers on collective action could simply assume a particular kind of interactive knowledge, and safely avoid discussing any messy details.

However, the details do matter. In this section, we review a particularly striking example in which only the highest form of interactive knowledge is sufficient for the agents to form the requisite collective intention. The ‘Coordinated Attack Problem’, which originated in the distributed systems literature (see Halpern 1986a), and which later attracted attention in the economics literature (see Rubinstein 1989; Binmore and Samuelson 2001), demonstrates a situation in which collective action is impossible because the necessary epistemic conditions cannot be met. The example is useful for the current philosophical project because it demonstrates important effects of interactive knowledge upon collective action. It thereby supports our contention that an account of interactive knowledge is necessary for a satisfactory account of collective action.

2.1 The structure of the Coordinated Attack Problem

It is frequently assumed that in a coordination problem, the agents have ‘common knowledge’ of the relevant features of the situation. For example, David Lewis’s seminal 1969 study of conventions makes this assumption. A standard interpretation⁴ of Lewis’s concept of common knowledge is that a proposition p is common knowledge if and only if:

(everyone knows that) ^{n} p

for all values of n . In other words, not only does everybody know that p , but everybody knows that everybody knows that p , and so on. Indeed, this common knowledge assumption, and the corresponding definition of common knowledge, have become standard. For example, it is standard to assume that the strategic features of a situation—including payoffs, risks, and so on—are ‘common knowledge’ in Lewis’s sense.⁵

Economists have long appreciated that the assumption of common knowledge is not innocuous. The so-called ‘Coordinated Attack Problem’⁶ is a striking example of this fact. It is a case in which common knowledge is required for action, but in which it is impossible for the agents to elevate a piece of information to the status of common knowledge. And this is in spite of the fact that the agents can communicate with each other indefinitely, have the same interests, and consciously aim to coordinate their beliefs and actions.

The example asks us to imagine that two allied armies are commanded by two generals (G_1 and G_2). In the valley below, the enemy is encamped. General G_1 knows that the enemy is below, and also knows that they will be victorious only if both generals attack at the same time. If either G_1 or G_2 attacks the enemy alone, the attacking general’s army will be defeated in a disastrous battle. On the other hand, if both attack simultaneously, they will win an easy victory. Thus, G_1 and G_2 face the problem of coordinating their attack. Communication between the two generals is through a messenger who can relay written messages between them. However, there is a small, but positive, probability ψ that the messenger will be unsuccessful in relaying any particular message from one general to another (he may be shot by the enemy or

⁴We say that this is a ‘standard interpretation’ instead of simply attributing this concept to Lewis because he does not formulate his concept of ‘common knowledge’ in exactly this way, although it is commonly represented in this way.

⁵For example, in Nash’s groundbreaking treatment of the equilibrium concept, the use of equilibria is justified, in part, by assuming that the agents have common knowledge of the game structure and common knowledge of each others’ rationality (see Nash 1950 and 1951).

⁶For our purposes here, the ‘Electronic Mail Game’ would serve equally well (see Halpern 1986a).

become lost, for example). We also assume that whenever a message is received by one of the generals, a confirmation message is automatically sent to the other general, and that confirmation message has the same small probability ψ of failing to reach its destination. It is an important feature of the example that the confirmation messages are sent automatically until one of them fails to reach its destination.

It is a surprising fact (proven in Rubinstein 1989 and Halpern 1986a) that, under these conditions, the generals cannot successfully coordinate their attack; this remains true no matter what finite number of confirmation messages are sent and received. We may show this by induction on the number of messages n received by G_2 . Suppose that $n = 0$. Then G_2 knows that either no messages were sent by G_1 , or that one message was sent by G_1 , but that it failed to get through. If the former is true, then it is clearly best for G_2 to not attack because G_1 has no intention of attacking. If the latter, then the situation is slightly more complicated. Because confirmation messages are sent automatically, G_2 knows that G_1 will not receive a confirmation message. Without such a confirmation message, G_1 will reason that one of the following two cases must be true:

- (1) G_1 's message did not reach G_2 . This has a probability of ψ .
- (2) G_1 's message reached G_2 , but G_2 's confirmation message did not reach G_1 . This has a probability of $\psi(1-\psi)$.

Because $\psi > \psi(1-\psi)$, then G_1 will believe that option (1) is more probable. But because G_2 will not attack if no message is received, then G_1 will not attack if G_1 believes that G_2 received no message. Therefore, G_1 will not attack. This fact is known to G_2 , so G_2 will not attack, either. This completes the basis case of the proof.

For the inductive step, we suppose that there is some n such that if n messages are received by G_2 , then neither G_1 nor G_2 will attack. Now suppose that exactly $n+1$ messages are received by G_2 . In this case, G_1 will not receive a confirmation message for message $(n+2)$, and this fact is known by G_2 . G_2 knows that G_1 will be faced with considering two options, which are closely analogous to (2) above, namely:

- (1) G_1 's first n messages reached G_2 , but the $(n+1)$ th message did not. This has a probability of $\psi(1-\psi)^n$.
- (2) G_1 's message reached G_2 , but G_2 's final confirmation message did not reach G_1 . This has a probability of $\psi(1-\psi)^{n+1}$.

At this point, G_2 believes that G_1 's line of reasoning will be the same as in the basis case. Option (1) is more probable because $\psi(1-\psi)^n > \psi(1-\psi)^{n+1}$, so G_2 believes that G_1 believes that G_2 probably did not receive the n^{th} message. But according to the inductive hypothesis, neither general will attack under that condition. This completes the inductive step. So no matter how many messages are sent and received, neither general will attack.

The conclusion that the generals cannot coordinate their attack is highly counterintuitive.⁷ But it does provide a rigorous demonstration of the most extreme possible case in which nothing short of full-blown common knowledge is sufficient for collective action.

On the other end of spectrum, it is easy to come up with cases in which a minimal level of interactive knowledge is necessary. For example, consider a simple case in which one choice is particularly salient. Two friends staying in the same hotel in a strange city get separated. It is obvious that they may be perfectly capable of meeting up together again at the hotel rather than, say, at a local landmark. The reason is that this choice is highly salient to both people. For our purposes here, we note that only a minimal level of interactive knowledge is necessary for them to achieve this coordination. Perhaps each needs only believe that the other person is capable of reaching the hotel, for example. No other peculiar features of the situation force them to acquire higher levels of interactive knowledge.

2.2 Issues raised by the Coordinated Attack Problem

By themselves, these two examples are sufficient to establish a negative thesis concerning the role of interactive knowledge in generating collective intentions. This is simply that there is no *unique* threshold level of interactive knowledge at which collective intentions are formed. Whether a particular level of interactive knowledge is sufficient depends upon the details of the situation in which the agents find themselves.

⁷In fact, it has been our experience that many philosophers of action are violently resistant to this conclusion. For those seeking a more rigorous demonstration, we refer the reader to Rubenstein's (1989) proof and Ken Binmore and Larry Samuelson's excellent 2001 discussion.

But in order to alleviate some of the discomfort surrounding this example, we point out that there are several assumptions in the Coordinated Attack Problem that are unrealistic. For example, it is necessary to assume that the generals can iterate their line of argument indefinitely, any finite number of times. It also assumes that the confirmation messages are sent automatically, any number of times. For the purposes of our discussion here, these assumptions are unimportant. We merely use the example to show that epistemic conditions for collective action may depend upon a variety of different factors, and that the epistemic conditions for collective action may be difficult to satisfy in some cases.

This observation thereby raises important questions for any account of collective intention. For example, if interactive knowledge is central to an account of collective intention, then we need to understand what features of the strategic situation determine the level of interactive knowledge required. Additionally, the paradoxical nature of the Coordinated Attack Problem raises the issue of when, if ever, full-blown common knowledge is required in realistic settings. And if common knowledge is ever a prerequisite for generating a collective intention, then we must ask how such knowledge becomes possible. Without answers to these questions, no account of collective intention (and thereby collective action) is complete.

3. Formalization of interactive knowledge

In order to address these questions, we require a framework for analysing interactive knowledge. Formal interactive epistemology contains both a syntax and a semantics for describing the effects of one agent's knowledge about another agent's knowledge. In what follows, we shall use the theory of interactive epistemology due to Aumann. Although there are some controversial assumptions implicit in Aumann's formalism, we shall argue that it provides a flexible and useful way of thinking about collective action and its epistemic conditions. And despite those assumptions, Aumann's work helps identify crucial assumptions that must be questioned on any satisfactory analysis of collective intention and action.

3.1 *Semantics of interactive knowledge*

In interactive epistemology, an agent's knowledge is identified as the set of states that she considers possible.⁸ Let Ω be the set of all the possible states of the world ω , and let $I = \{1, 2, \dots, n\}$ be a set of agents. Events are identified with subsets of Ω . For example, if Ω consists of all the possible ways that the weather might be tomorrow, then the event of its raining tomorrow is identified with the set consisting of all the elements $\omega \in \Omega$ that are states in which it rains tomorrow. For each individual agent $i \in I$, we make the following observations and definitions in terms of the set of states Ω .

First, we allow that the agents might be unable to differentiate among some of the elements of Ω . For example, we might imagine a situation in which a die is cast, and an agent is told only whether the die has landed on an even or an odd number. In this case, we define the set of

⁸For excellent surveys, see Geanakoplos 1992 and Koessler 2000.

possible rolls of the die $\Omega = \{1, 2, \dots, 6\}$. Suppose that the die is cast and lands on 4. Then the actual state of the world is $\omega = 4$. But since the agent i is told only that the die landed on an even number, then i knows only that the real state of the world is some $\omega' \in \{2, 4, 6\}$. We represent this incomplete knowledge by defining agent i 's *knowledge function* $\kappa_i: \Omega \rightarrow \mathcal{P}(\Omega)$, where $\mathcal{P}(\Omega)$ is the set of all subsets of Ω . If the true state of the world is ω , then $\kappa_i(\omega)$ is the set of all $\omega' \in \Omega$ that i considers possible. So in our example, $\kappa_i(4) = \{2, 4, 6\}$.

The function κ_i induces a partitioning of Ω into subsets of worlds between which i cannot differentiate. In the above example, the agent cannot differentiate between states 2, 4, and 6 (since she is told only whether the die has landed on an even or an odd number). Similarly, she cannot differentiate between 1, 3, and 5. So κ_i partitions Ω into $\{1, 3, 5\}$ and $\{2, 4, 6\}$. This set of partitions is defined by the function I_i as follows:

$$(1) \quad I_i(\omega) = \{\omega' \in \Omega \mid \kappa_i(\omega) = \kappa_i(\omega')\}$$

It is helpful to single out the partition of possible worlds that i considers possible, given the actual state of the world. This partition is i 's *information partition*, and is designated I_i .

As we have said, an event E is identified with the set of possible worlds in which that event happens. Accordingly, Aumann's theory of interactive epistemology makes use of the following definition of an agent's knowledge:

Definition 1. An agent i knows an event E if and only if every world i considers possible is a (possibly improper) subset of E . Formally, i knows that E if and only if $\kappa_i(\omega) \subseteq E$, where ω is the actual world.

In order to employ this account for interactive knowledge (that is, for situations in which i knows that j knows that ...), an agent's knowing an event must itself be an event, so that this epistemic state can itself be the object of an agent's knowledge. So the following definition is required:

Definition 2. The event of i 's knowing that E , denoted $K_i E$, is the set of all worlds that i considers possible, and where the event E obtains. Formally, $K_i E = \{\omega \mid \omega \in I_i \text{ and } \omega \in E\}$

We now show how interactive knowledge is represented in this theory.

Theorem 1. For agents i, j , and an event E , i knows that j knows that E if and only if $\kappa_j(\kappa_i(\omega)) \subseteq E$.

Proof. We follow the proof in Geanakoplos 1992. First, we prove the ‘only if’ direction. Suppose i knows that j knows that E . Now assume for *reductio* that there is some world $\omega' \in \kappa_i(\omega)$ which i considers possible such that $\kappa_j(\omega') \subseteq E$. But if so, then ‘from the perspective’ of i , ω' has the following two properties:

- (1) ω' is a world i considers possible, and
- (2) if ω' were the actual world, then j 's knowledge function would have it be the case that E might be false, and hence, E could not be known to j . But this is contrary to the assumption that j knows that E .

This suffices to show the first direction. We now show the ‘if’ direction (by contraposition). Suppose that i does not know that j knows that E . Then there is some $\omega' \in \kappa_i(\omega)$ such that $\kappa_j(\omega') \subseteq E$ is false. But then, it is also false that $\kappa_j(\kappa_i(\omega)) \subseteq E$.

Since Theorem (1) gives us a method for representing higher-order knowledge, we now rehearse some standard definitions.

Definition 3. An event E is *mutual knowledge* among agents $I = \{1, 2, \dots, n\}$ if and only if $\kappa_i(\omega) \subseteq E$, for all $i \in I$. E is *known to level k* among the members of I if and only if for all $m \leq k$, and every series x_1, x_2, \dots, x_m with each $x_i \in I$, it is true that $\kappa_{x_1}(\kappa_{x_2}(\dots \kappa_{x_m}(\omega) \dots)) \subseteq E$. E is *common knowledge* among the members of I if and only if E is known to level m , for all $m \in \mathbb{N}$. E is *almost common knowledge* if E is known to level k , and k is large but finite.

3.2 Interactive knowledge theses

We may now state the following (appropriately controversial) theses about knowledge.

- (A1) $K_i(E \cap F) = K_i E \cap K_i F$ (Distribution Axiom) An agent knows a conjunction if and only if she knows both conjuncts.
- (A2) $K_i E \subseteq E$ (Truth Axiom) If an agent knows a proposition, then that proposition is true.
- (A3) $K_i E \subseteq K_i^2 E$ (Introspection Axiom) If an agent knows a proposition, then she knows that she knows it.
- (A4) $K_i E \subseteq K_i K_i E$ (Wisdom Axiom) If an agent does not know a proposition, then she knows that she does not know it.

In the interests of full disclosure, we note here that theses (A1)–(A4), although highly dubious, are validated by any semantics of common knowledge that uses information partitions in the way we have outlined above. However, when we apply this framework below, our arguments will not depend upon any of them.

Theorem 2. Theses (A1)–(A4) are validated by the knowledge semantics above.

Proof. We show that (A3) is valid. Suppose that $K_i E$, but that $K_i K_i E$ is false. So $K_i E = \{\omega \mid \kappa_i(\omega) \text{ and } \omega \in E\}$. But by the same token, $K_i K_i E = \{\omega \mid \kappa_i(\kappa_i(\omega)) \text{ and } \omega \in \kappa_i(E)\}$. But clearly, the latter is a (possibly improper) subset of the former. The remaining theses follow analogously.

4. Applications of interactive knowledge

Aumann's theory is valuable because it enables us to inquire more precisely into the epistemic requirements for collective action. Those epistemic requirements, so we shall begin to argue in this section, illuminate several important issues in philosophical action theory.

4.1 Specifying the level of interactive knowledge

In order to analyse the concept of collective action, we must specify what distinguishes mere collections of actions from a single collective action. It is clear from any reasonable set of examples that an important distinction between mere collections of actions and collective actions will make recourse to the beliefs that agents have about the beliefs of the other agents in the collective. For example, if several people happen to show up at a restaurant at the same time, they do not do so because they have any particular beliefs about the other people at the restaurant. But on the other hand, when a group deliberately meets at a restaurant in order to have lunch together, each individual shows up at the appointed time because each believes of the others that they believe that it is the appointed place and time. In such cases, the presence of interactive knowledge is an important mark of collective action.

While interactive knowledge marks this important distinction, some writers on collective action have left the notion as an unanalysed assumption in their analyses. For example, Michael Bratman (1993, p. 99) notes the importance of interactive knowledge, but specifically leaves aside the issue of how to characterize it.

Rather than leaving it as an undefined primitive, other writers specify exactly what level of interactive knowledge is required for the performance of a collective action. For example, in his discussion of ‘joint intentions’ (which we have been calling more generically ‘collective intentions’), Tuomela appears to specify a particular level of interactive knowledge. In discussing his taxonomy of intentions, he argues that agents have a joint intention if and only if each one has a particular kind of individual intention, which he calls a ‘we-intention’, and those we-intentions are related to each other in a particular way. His criterion for when a group of agents has a joint intention is the following:

Agents $A_1, \dots, A_i, \dots, A_m$ have the joint intention to perform a joint action X if and only if:

- (a) these agents have the we-intention (or are disposed to form the we-intention) to perform X ; and
- (b) there is a mutual belief among them to the effect that (a).
(Tuomela 2005, p. 342)

Because of the specificity of condition (b), Tuomela’s analysis implies the existence of a specific threshold level of interactive knowledge that is necessary and sufficient for satisfying the epistemic conditions for the formation of a so-called ‘joint intention’.

Another commonly employed strategy is to require full-blown common knowledge (in Lewis’s sense) for the performance of a collective action, while leaving open the possibility that there may be cases in which something short of this high level of interactive knowledge is required. Here, the motivating idea seems to be that if some level of interactive knowledge is required for collective action, then certainly full-blown common knowledge would also be sufficient—after all, if an agent has common knowledge, then it is logically entailed that they also possess every weaker form of interactive knowledge as well. Christopher Kutz is an example of a writer whose theory falls into this category, since (as we have seen above) he assumes that higher levels of interactive knowledge will facilitate collective action.

4.1.1 *When (interactive) ignorance is bliss*

However, all three of these approaches have serious drawbacks. Consider the third approach. Although it seems reasonable to suppose that common knowledge will facilitate collective action, this is not always the case. Consider the following example. Two agents— A and B —are to decide whether to make a wager, which is depicted in Figure 1.

The wager is to guess whether the point p is located inside the dashed lines; player A wins \$1 from player B if point p is inside the lines, and player B wins \$1 from A otherwise. The nine squares in Figure 1 represent the information partitions for both players (so we assume that $\Omega_A = \Omega_B$). Recall that an information partition contains states of affairs among which an agent cannot differentiate. Thus, even if the players knew the square precisely in which point p is located, neither would know the exact location of p . So we assume that a player will accept the bet if and only if she judges it more likely than not that she will win. After the bet has been placed, the precise location of point p is revealed, and the winner collects her payoff from the loser.




Figure 1: A wager in which Player A has knowledge of the column of point p , and Player B has knowledge of the row of point p

Initially, Player A knows only that p is within column a_3 , and Player B knows only that p is within row b_1 . Accordingly, both A and B are willing to accept the bet—for most points in a_3 are inside the dashed figure (inducing A to accept the bet), while most points inside b_1 are not (which induces B to accept). But if they repeatedly learn that the other player is willing to accept the bet, then the situation changes. For suppose A announces her willingness to accept the bet. Then B knows that A knows that p is not in the region a_1 (or else A would not have been willing to accept the bet). But still, B is willing to accept, since she now knows that point p must be within either (b_1, a_2) or (b_1, a_3) . Most points in (b_1, a_2) and (b_1, a_3) are outside the dashed lines, so B is willing to accept the bet. Suppose that B announces her willingness to accept. Now A knows that B knows that A is willing to accept the bet, and A knows that B's knowledge does not prevent her from accepting. So A

can reason that p cannot be in b_2 or b_3 . So now, A knows that point p is in (b_1, a_3) . But when A announces his continued willingness to accept the bet, B will infer that the point must be inside (b_1, a_3) . And because this information would give the advantage to Player A , Player B will call off the bet. Thus, in spite of the fact that both players were initially willing to accept the bet, the bet is now off.

In this example, each announcement by a player creates higher-order knowledge. Initially, the players have different information sets, and lack even mutual knowledge about the probable location of point p . But as they repeatedly announce their willingness to accept the wager, they generate higher-order knowledge and their information sets gradually converge. As soon as their knowledge is close enough to common knowledge given the coarseness of Ω , they are unable to agree on the wager. Therefore, we have an example in which a collective action (making a wager) can take place only if the agents *lack* higher-order knowledge.

Furthermore, the example can be straightforwardly modified so that higher-order interactive knowledge *facilitates* collective action. Suppose that the players are no longer making a wager against each other. Instead, let us suppose that players A and B are both betting against a third party C . If both A and B accept the offer, then they will each win \$1 if point p is outside the dashed lines. If they both accept and point p is inside the dashed lines, then each has to pay C \$1. But if either A or B is unwilling to accept the wager, then the bet is off.

As before, suppose that A initially knows only that p is within a_3 , while B knows only that p is within b_1 . Initially, therefore, A is unwilling to accept the wager, while B is. But if they repeatedly announce their willingness or unwillingness to accept the bet with C , then eventually both A and B will be willing to make the wager. Therefore, this is a case in which higher-order knowledge *facilitates* collective action.

At this point, we need to address a possible worry regarding our use of a competitive activity—namely, gambling—as a form of cooperative behavior. After all, it is reasonable to see activities such as gambling as ones that are the opposite of collective actions. Indeed, one might be tempted to say that any activity in which the individuals' interests conflict cannot be a collective action in the relevant sense. To address this issue, it is useful to consider a discussion from John Searle:

It is worth noticing in passing that most forms of competitive and aggressive behavior are forms of higher level cooperation. Two men engaged in a prize fight are engaged in a form of competition, but it is a form of aggressive competition that exists only within a higher level form of cooperation. Each prize

fighter has the intention to hurt the other, but they have these intentions only within the frame of the higher order intention to cooperate with each other in engaging in a prize fight. This is the distinction between a prize fight and a case of one man simply assaulting another man in a dark alley. And what goes for the prize fight also goes for football games, business competitions, courtroom trials, and in many cases, even armed warfare. For human beings, most social forms of aggressive behavior require higher level cooperation. (Searle 1990, p. 20)

In the above passage, Searle is pointing out that even in activities in which the individuals' interests conflict, there is often a higher-order set of rules to which the agents conform. In such cases, the higher-order rules may be a precondition for the competitive activity.

To take Searle's example, if two prize fighters do not cooperate in respecting the rules of the competition (e.g. by returning to their respective corners when the bell rings), then they cannot engage in the fight. Similarly, two individuals cannot gamble if they do not cooperate in conforming to the (perhaps unspoken) rules of gambling. So if one were to focus on the competitive, lower-level aspects of the activity, one would miss the equally important cooperative preconditions of the activity.

Although it may be surprising that there are cases in which higher-order knowledge either blocks or facilitates collective action, this example actually illustrates a very general phenomenon—and one that is illuminated by Aumann's theory of interactive knowledge. In a result that is well-known in economics, but which is underappreciated by philosophers concerned with collective action, Aumann has shown the following:

Theorem 3. Suppose two agents A and B assign common prior probabilities for an event E . Suppose that after receiving (possibly different) information about the probability of E , and that after receiving this information, the probabilities that they assign to E are $P_A(E)$ and $P_B(E)$. If they have common knowledge of these subjective probabilities, then $P_A(E) = P_B(E)$.

Proof. See Aumann 1999b.

Theorem (3) is relevant to theories of collective action because it shows that when two or more rational agents have common knowledge of each others' subjective probability assignments, then they will necessarily agree on those probabilities. Hence, if the two agents are in a situation where such agreement facilitates collective action, then common knowledge will also facilitate collective action. But as we have seen in the gambling example, it is easy to come up with cases in which collective action

is blocked when the agents agree on the probability assessments. In cases of this type, common knowledge of each others' subjective probability assignments will block collective action. In this way, Theorem (3) gives us a general reason to be sceptical of any account of collective action, which, like the theory offered by Kutz, assumes that high levels of interactive knowledge *always* facilitate collective action.

4.1.2 When is common knowledge appropriate?

So far, we have learned two lessons. First, the Coordinated Attack Problem shows that it is not always possible for two agents to reach a state of common knowledge, even when it is in their best interests to perform a collective action, and even when they are able to achieve any finite level of interactive knowledge. Second, there are cases in which common knowledge is deleterious to collective action. This is the lesson we learn from the gambling example depicted in Figure 1. Thus, it is reasonable to wonder under what conditions rational agents will be able to acquire common knowledge. Here, too, interactive epistemology is useful.

An example will help to make this clear. Consider the game depicted in Figure 2, which we shall call the 'Risk Game'.

		Player B		
		<i>a</i>	<i>b</i>	<i>c</i>
Player A	<i>a</i>	5,5	-5,-5	-5,0
	<i>b</i>	-5,-5	5,5	-5,0
	<i>c</i>	0,-5	0,-5	0,0

Figure 2: The Risk Game: A game in which some level of interactive knowledge is required if risk-averse players are to coordinate their actions

In this situation, the players get a positive payoff for coordinating their actions on either *a* or *b*. But if a player chooses either *a* or *b*, and the other player does not match her choice, then the player is penalized by receiving a -5 payoff. However, the option *c* is available to any player as an 'outside option'. If a player chooses option *c*, then she receives the *status quo* payoff of 0, regardless of the choice of her partner.

In a Risk Game, it is not enough for the players to conditionalize their actions upon some outside signal, as in the Coordination Game. For a risk-averse player may judge that even though she received the signal, her partner might not have received it. In that situation, she

risks receiving the -5 penalty. Importantly, this is true even if both players successfully receive the signal, because each player will play either a or b only if she knows that the other player also received the signal. Thus, mutual knowledge may be insufficient for coordination in a Risk Game, especially if the risk is very great.

Of course, if both players were to coordinate their actions on (say) choice a , then it would be natural to interpret their joint behavior as a collective action. In fact, we may fairly consider the Coordinated Attack problem to be an extreme case of a Risk Game, and we judge that, had the generals coordinated their attack, then they would thereby have performed a collective action.

Indeed, it is often the role of public signals to establish common knowledge. To illustrate, note that one of the paradoxical features of the Coordinated Attack Problem is that the two generals are unable to establish common knowledge in spite of their sophisticated message schemes; in contrast, there are many cases in which it is easy for two agents to establish common knowledge. For example, suppose that the two generals are on high hilltops, and that one general signals to the other by raising a flag. Suppose further that each knows that anyone on a hilltop can see the flag, and that they both see each other standing on their respective hilltops. In this sort of situation, they have all the beliefs necessary to answer affirmatively any question of the form, 'does the other General know that I know that ... observed the flag?' Thus, they have common knowledge that the flag was raised.⁹

In cases in which there is a public signal, and it is possible for all the agents to observe each other observing the signal, common knowledge may be easy to come by. This either aids or hinders the performance of the collective action, depending on the structure of the underlying strategic situation. However, there is another option that applies to more cases—this is when the agents display bounded rationality.

It is a peculiar feature of the Coordinated Attack Problem that for any $n \neq m$, the agents are able to differentiate between n th-order and m th-order knowledge. Although this may be a reasonable assumption so long as m and n are very small, it is clearly too much to require that (say) an agent calculates based upon fifty-seventh-order knowledge, and distinguishes that knowledge from fifty-eighth-order knowledge.¹⁰

⁹ For a detailed explanation of how common knowledge arises in such a situation, see Lewis 1969, Ch. 1, in which Lewis details the argument that each individual would use when reasoning about each other's knowledge.

¹⁰ It is therefore an important empirical question whether human beings frequently differentiate between n th-order and $n+1$ th-order knowledge, for large values of n . Colin Camerer has taken up this question in his 2003.

However, even if agents are unable to calculate beyond a small, finite degree of interactive knowledge, we are confronted with the question of whether that limitation facilitates or blocks any particular instance of collective action. Surprisingly, under certain conditions, this limitation transforms higher-order knowledge into common knowledge, and thereby may facilitate collective action. In particular, we have the following result.

Theorem 4. Let G be a group of agents, all of whom have a finite number of information partitions. There exists a finite number n such that if the members of G have interactive knowledge to level n of an event E , then E is common knowledge among G .

Proof. Consider a group G of individuals, and assume that for all $i, j \in G$ $\Omega_i = \Omega_j$ so that we may simply refer to any individual's information partition as Ω . It is easy to show that this assumption is harmless, although dropping the assumption makes the proof cumbersome.

Define the closure of ω for G , denoted $CL(\omega, G)$, as the smallest set meeting the following two conditions:

1. $\omega \in CL(\omega, G)$
2. For any $i \in G$, if $\omega' \in CL(\omega, G)$ then $\kappa_i(\omega') \in CL(\omega, G)$

Clearly, if Ω is finite, then so is $CL(\omega, G)$.

Consider an event E and an arbitrary $i \in G$. In order for E to be known to level k , then for every $m \leq k$, and every $x_i \in G$, $\kappa_{x_i}(\kappa_{x_2}(\dots \kappa_{x_m}(\omega) \dots)) \subseteq E$. Clearly, there are at most $\sum_{x=0}^{k-1} |G|^x$ statements that must be true.

Choose a value of k such that $\sum_{x=0}^{k-1} |G|^x > |CL(\omega, G)|$. We are guaranteed that there is such a value of k because $CL(\omega, G)$ is finite if Ω is. Suppose that E is known to level k . It is sufficient to show that E is also known to level $k+1$. For this, it is sufficient to show that $\kappa_{x_1}(\kappa_{x_2}(\dots \kappa_{x_{k+1}}(\omega) \dots)) \subseteq E$ for some arbitrary series of x_i s in G .

Because $\sum_{x=0}^{k-1} |G|^x > |CL(\omega, G)|$, then by a pigeonhole argument, there must be some $\kappa_{y_1}(\kappa_{y_2}(\dots \kappa_{y_\ell}(\omega) \dots))$, $\ell < k+1$ such that $\kappa_{y_1}(\kappa_{y_2}(\dots \kappa_{y_\ell}(\omega) \dots)) = \kappa_{x_1}(\kappa_{x_2}(\dots \kappa_{x_{k+1}}(\omega) \dots))$. But by assumption, $\kappa_{x_1}(\kappa_{x_2}(\dots \kappa_{x_{k+1}}(\omega) \dots)) \subseteq E$, so $\kappa_{y_1}(\kappa_{y_2}(\dots \kappa_{y_\ell}(\omega) \dots)) \subseteq E$, which completes the proof.

Theorem 4 has a simple informal justification. Recall that two possible worlds are elements of the same partition of Ω if and only if the agent is unable to distinguish between them. So agents who often cannot distinguish between various possible worlds will have partitions that place many different possible worlds in the same equivalence class of Ω . The condition that $|\Omega|$ is finite therefore amounts to saying that the agents can distinguish among only a finite number of possible worlds. In this case, there will be some highest order of interactive knowledge above which the agent cannot discriminate. Thus, any set of agents achieving that level of interactive knowledge with respect to some proposition is in an epistemic state that is indistinguishable (given their cognitive limitations) from common knowledge.

It is also helpful to informally motivate how the value of n is calculated in Theorem 4. The theorem assures us that for each agent A , there is some finite number n such that if A has interactive knowledge of some proposition p to degree n , then A 's interactive knowledge of p is indistinguishable from full-blown common knowledge. Let us call n the 'threshold' for A . The question is what Theorem 4 tells us about how A 's threshold is determined.

For simplicity, let us consider a simple case in which there are only two agents, whom we shall call A and B . The theory of interactive epistemology using information partitions identifies each event—including the event of some proposition being known by an agent—with some information partition. So consider a unary function ' \square ' over information partitions, where ' $\square p$ ' is interpreted as ' A knows that B knows that p '. Because the semantics of interactive knowledge treats the knowledge of an event as just another event, the \square function will take information partitions and map them onto information partitions. Theorem 4 assures us that if there are only finitely many information partitions, then there will be some threshold n such that $\square^n p$ is the same partition as $\square^m p$, where $m < n$. A trivial induction shows that for any $\ell > n$, there will exist some $\ell' \leq n$ such that $\square^\ell p = \square^{\ell'} p$. So if A 's interactive knowledge of p has already reached the threshold n , then A has every higher level of interactive knowledge of p also.

The existence of some finite threshold value n is guaranteed if $|\Omega|$ is finite. But it is important to note that the existence of such a threshold does not require $|\Omega|$ to be finite. All that is required is that there be some $n > m$ such that $\square^n p = \square^m p$. Of course, because Theorem (4) is only a formal result, it is appropriately silent about any empirical facts that would explain *why* some higher level of interactive knowledge would be treated like a lower level of interactive knowledge. But an explanation

for many such cases would be that the agents have cognitive limitations that prevent them from distinguishing among every possible level of interactive knowledge. This simple observation sheds some light on the paradoxical flavour of the Coordinated Attack Problem. Specifically, the two generals are hurt by their own cognitive superiority. Because each is assumed to be able to distinguish between any two distinct levels of higher-order knowledge, then they must have (at least) a countably infinite partition of Ω . If their rationality were bounded, then after some finite number of confirmations, we would expect them to successfully coordinate their actions. And of course, this is just what we would intuitively expect.

Because we have argued that an agent's having lower levels of interactive knowledge may *sometimes* entail her having higher levels of interactive knowledge, we raise the question of whether this is *generally* the case. In particular, if an agent's threshold value is n , then the question is whether, having reached interactive knowledge of level $n-1$, the agent will be able to reason as if she had interactive knowledge to degree n , and thereby achieve the equivalent of common knowledge.

We believe that there is no univocal answer to this question. Suppose that an agent A reasons to n th-order interactive knowledge of p . It may turn out that the reasoning employed by A may in principle be reiterated without additional information to level $n+1$. Thus, if A is rational enough to recognize that fact and carry out the induction, then A 's threshold may, in effect, be lowered. But from this example, we may *not* draw the conclusion that cognitive superiority and idealized rationality must allow an agent to bootstrap her interactive knowledge to degree n into degree $n+1$. Indeed, this is the lesson we learn from the Coordinated Attack Problem, in which we have ideally rational agents who are prevented from leveraging arbitrarily high levels of interactive knowledge into common knowledge.

4.2 Mutual knowledge

In spite of the emphasis on interactive knowledge, there are also cases in which no interactive knowledge is required for the performance of collective action. Recall that a proposition p is mutual knowledge if everyone knows that p ; mutual knowledge does not require that the agents are aware of each others' knowledge of p .

Consider the simple Coordination Game depicted in Figure 3. In it, players receive a payoff if they coordinate their actions, but receive nothing if they miscoordinate. Here, only mutual knowledge is required because the penalty for miscoordination is no worse than the

status quo (for we assume that if the players simply walked away from the game, then each would receive a payoff of zero), and there is no incentive for either party to deliberately miscoordinate.

		Player II	
		<i>a</i>	<i>b</i>
Player I	<i>a</i>	1,1	0,0
	<i>b</i>	0,0	1,1

Figure 3: The Coordination Game

For example, suppose that you and I are playing a coordination game, and we have decided that we shall conditionalize our behavior on some external signal; let us assume that if we observe the signal, then we will play *A*, otherwise we will play *B*. Having determined this strategy, suppose that you observe the signal. Clearly, it is in your best interest to play *A*, regardless of whether you know that I have observed the signal. After all, the structure of the game does not give me any incentive to deliberately miscoordinate, or ‘play it safe’ by not playing strategy *A*. Because the situation is symmetric, I shall also reason in the same way. So if we have mutual knowledge of the signal, then that weak level of interactive knowledge will be sufficient to ensure that we perform the collective action of coordinating our strategies.

5. Solving the problem of Alfonse and Gaston

David Velleman has directed a particularly interesting objection to Michael Bratman’s theory of shared cooperative activity; a response to this objection is suggested by considering the role played by interactive knowledge.

Velleman (1997, p. 35) argues that Bratman’s theory of collective intention would fail to ever produce a collective action because each individual’s intention is conditional upon the intentions of the other person. For example, Bratman discusses a case in which two people are required in order to pump water into a house—let us say that Abe must pump the water while Betty turns a valve. On Bratman’s account of collective intention, Abe will have the intention to pump the water if Betty turns the valve, while Betty will have the intention to turn the valve if Abe pumps the water. According to Velleman, these two ‘conditional’ intentions would be insufficient to bring about their collective action. Instead, they will behave like the vaudeville characters Alfonse and Gas-

ton, who are unable to walk through a doorway because each of them politely waits for the other to go first.

Bratman's response to this worry is to point out that in such a case, one person's behavior can be settled by observing that the other has 'done his part'. In the example, we may suppose that Abe recognizes Betty's conditional intention, and therefore simply starts pumping the water, knowing that Betty will observe him doing so and will thereby be encouraged to turn the valve. Here, Bratman is correct to point out that either person could rationally do her part of the collective action, acting on the assumption that by doing so, the other member of the group will be led to perform her part of the collective action. Bratman summarizes his response in the following way:

... I can 'frame' the intention that we *J* in part on the assumption that you will, as a result, come also so to intend. While I confidently predict you will come so to intend, I also recognize that you remain a free agent and this decision is really up to you ... Second, even after I have formed the intention that we *J*, in part because I predict you will concur, I can recognize that you still need to concur: It is just that I am fully confident that you will. (Bratman 1999, p. 157)

Bratman's discussion of this particular case is perfectly sound; indeed, it is easy to come up with similar examples in which one person's intention to participate in the collective action is settled when that person observes the participation of others. However, we believe that considerations of interactive knowledge provide a more general response to this kind of worry, and we also believe that a more general response is necessary. After all, it is important to note that an 'Alfonse and Gaston' objection could be made in regard to cases where no person in the group can observe the others' participation. A simple example is the one discussed above, in which we must coordinate our actions contingent upon some external signal. It may be the case that we are perfectly justified in adjusting our strategies accordingly, even if we are unable to observe each other. In such a case, Bratman's response will not apply.

However, an interactive knowledge framework makes it clear that Bratman's example is simply one way in which an agent can gain interactive knowledge of another agent's beliefs and actions. In other words, what makes Bratman's case a reasonable reply to Velleman's objection is that the case demonstrates that the presence of interactive knowledge allows the agents to avoid becoming mired in a situation where their conditional intentions are never executed. For when the level of interactive knowledge is appropriate to the strategic situation, it may become rational for each agent to act.

Once we learn this general lesson from Bratman's example, it becomes clear that there are many other ways in which an appropriate level of interactive knowledge can be achieved. And indeed, as we have seen above, there are cases in which the level of interactive knowledge required is extremely low—so low, in fact, that the agents may rationally act without having any assurance that the other agents will act similarly.

6. Conclusion

In this paper, we have argued that theories of collective action can benefit by taking into account the role played by interactive knowledge. In one respect, this lesson is not new—for as we have seen, virtually all accounts of collective intention make recourse to interactive knowledge in some way or other. However, the obviousness of the fact that interactive knowledge does play such a role belies the fact that it may do so in complex ways.

Accordingly, we have argued in this paper that the role played by interactive knowledge is much richer and more complex than has been appreciated. In particular, different levels of interactive knowledge may be required, depending upon the strategic situation in which the agents find themselves. Furthermore, interactive knowledge is not always helpful in bringing about collective action—occasionally, collective action will be possible only when a particular threshold level of interactive knowledge is *not* reached.

Interactive knowledge has an unusual place in the literature on collective intention. On the one hand, interactive knowledge is taken to be of central importance to virtually all accounts of collective intention. But surprisingly, it is not an exaggeration to say that almost no effort has been devoted to analysing it. Thus, we believe that the potential payoff for analysing interactive knowledge—particularly within Aumann's framework—is substantial. For we suspect that a number of open problems that arise in discussions of collective intention and collective action may be illuminated by an account of interactive knowledge. To take just one example, if moral responsibility depends upon the epistemic state of the agent—as seems perfectly reasonable to suppose—then it may turn out that analyses of collective moral responsibility can benefit from a better understanding of how interactive knowledge is distributed and generated within a group.

Department of Philosophy
University of Missouri-Columbia
438 Strickland Hall
Columbia
Missouri 65211
USA

SARA RACHEL CHANT
AND ZACHARY ERNST

References

- Aumann, Robert J. 1999a: 'Interactive Epistemology I: Knowledge'. *International Journal of Game Theory*, 28, pp. 263–300.
- Aumann, Robert J. 1999b: 'Interactive Epistemology II: Probability'. *International Journal of Game Theory*, 28, pp. 301–14.
- Binmore, Ken and Larry Samuelson 2001: 'Coordinated Action in the Electronic Mail Game'. *Games and Economic Behavior*, 35, pp. 6–30.
- Bratman, Michael E. 1992: 'Shared Cooperative Activity'. *The Philosophical Review*, 101 (2), pp. 327–41.
- Bratman, Michael E. 1993: 'Shared Intention'. *Ethics*, 104, pp. 97–113.
- Bratman, Michael E. 1999: *Faces of Intention: Selected Essays on Intention and Agency* (Cambridge Studies in Philosophy). Cambridge: Cambridge University Press.
- Camerer, Colin 2003: *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.
- Cohen, Philip R. and Jerry Morgan (eds) 1990: *Intentions in Communication*. Cambridge, MA: MIT Press.
- Geanakoplos, John 1992: 'Common Knowledge'. *The Journal of Economic Perspectives*, 6(4), pp. 53–82.
- Gilbert, Margaret 1990: 'Walking Together: a Paradigmatic Social Phenomenon'. *Midwest Studies in Philosophy*, 15, pp. 1–14.
- Halpern, Joseph Y. 1986a: 'Reasoning About Knowledge: An Overview'. In Halpern 1986b, pp. 1–18.
- Halpern, Joseph Y. (ed.) 1986b: *Reasoning about Knowledge: Proceedings of the 1986 Conference*. San Francisco, CA: Morgan Kaufmann.
- Koessler, Frederic 2000: 'Common Knowledge and Interactive Behaviors: A Survey'. *European Journal of Economic and Social Systems*, 14(3), pp. 271–308.
- Kutz, Christopher 2000: 'Acting Together'. *Philosophy and Phenomenological Research*, 61(1), pp. 1–31.
- Lewis, David 1969: *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.
- Nash, John 1950: 'The Bargaining Problem'. *Econometrica*, 18, pp. 155–62.

- Nash, John 1951: 'Non-Cooperative Games'. *Annals of Mathematics*, 54(2), pp. 286–95.
- Rubinstein, Ariel 1989: 'The Electronic Mail Game: Strategic Behavior under "Almost Common Knowledge"'. *American Economic Review*, 79, pp. 385–91.
- Searle, John 1990: 'Collective Intentions and Actions'. In Cohen and Morgan 1990, pp. 401–16.
- Searle, John 1995: *The Construction of Social Reality*. New York, NY: The Free Press.
- Searle, John 1997: 'Responses to Critics'. *Philosophy and Phenomenological Research*, 57(2), pp. 449–58.
- Tuomela, Raimo 1984: *A Theory of Social Action*. Dordrecht: D. Reidel Publishing Company.
- Tuomela, Raimo 1989: 'Actions by Collectives'. *Philosophical Perspectives*, 3, pp. 471–96.
- Tuomela, Raimo 1991: 'We Will Do It: An Analysis of Group-Intentions'. *Philosophy and Phenomenological Research*, 51(2), pp. 249–77.
- Tuomela, Raimo 2002: *The Philosophy of Social Practices: A Collective Acceptance View*. Cambridge: Cambridge University Press.
- Tuomela, Raimo 2005: 'We-Intentions Revisited'. *Philosophical Studies*, 125, pp. 327–69.
- Velleman, J. David 1997: 'How to Share an Intention'. *Philosophy and Phenomenological Research*, 57(1), pp. 29–50.