

[Jonathan Birch](#)

Joint know-how

**Article (Published version)
(Refereed)**

Original citation:

Birch, Jonathan (2018) *Joint know-how*. [Philosophical Studies](#). ISSN 0031-8116

DOI: [10.1007/s11098-018-1176-6](https://doi.org/10.1007/s11098-018-1176-6)

Reuse of this item is permitted through licensing under the Creative Commons:

© 2018 The Author
CC BY 4.0

This version available at: <http://eprints.lse.ac.uk/90522/>

Available in LSE Research Online: October 2018

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

Joint know-how

Jonathan Birch¹ 

© The Author(s) 2018

Abstract When two agents engage in a joint action, such as rowing together, they exercise joint know-how. But what is the relationship between the joint know-how of the two agents and the know-how each agent possesses individually? I construct an “active mutual enablement” (AME) account of this relationship, according to which joint know-how arises when each agent knows how to predict, monitor, and make failure-averting adjustments in response to the behaviour of the other agent, while actively enabling the other to make such adjustments. I defend the AME account from three objections, and I then use this account as the platform for an examination of the reducibility (or otherwise) of joint know-how to joint propositional knowledge. A summative account of joint propositional knowledge is incompatible with the reduction of joint know-how to joint propositional knowledge, whereas a distributive account is not (although serious difficulties for any such reduction remain). I close by highlighting some open questions the AME account brings into view concerning the evolutionary origin and scaling up of joint know-how.

Keywords Joint action · Joint intention · Coordination · Cooperation · Know-how · Intellectualism

✉ Jonathan Birch
j.birch2@lse.ac.uk

¹ Department of Philosophy, Logic and Scientific Method, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK

1 Introduction

Collaboration in pursuit of joint goals is a ubiquitous aspect of human lives, and it was fundamental to the social foraging of our Palaeolithic ancestors. Yet it seems absent from the social lives of our fellow great apes. As the primatologist Michael Tomasello has remarked, “it is inconceivable that you would ever see two chimpanzees carrying a log together”.¹ Chimpanzees cooperate, but their cooperation takes the form of the pursuit of individual goals that happen to align. For example, two or more chimpanzees may hunt a monkey at the same time, but each is merely pursuing the prey individually: there is, Tomasello argues, no joint goal in these cases, and no conceiving of individual actions as means to that goal (Tomasello 2014).

This points to a central role for joint action in the story of human evolution (Sterelny 2012; Tomasello 2014, 2016). For this reason and others, recent decades have seen growing interest across various disciplines in the basic psychological structures underlying joint action, often under the broad heading of joint or shared intentionality. Most work in this area falls into one of two traditions. First, there are attempts to capture, at the personal level, the nature of the propositional attitudes possessed by the agents in a joint action, and in particular the nature of the intentions they execute. A landmark in this area is Michael Bratman’s (1992, 1993, 2014) theory of shared intention, on which, roughly speaking, a shared intention consists of a package of individual intentions which interlock in a characteristic way.² Second, there are attempts to capture, at the subpersonal level, the cognitive and neural mechanisms implicated in the skilful control of a joint action. A landmark in this area is Cordula Vesper and colleagues’ minimal cognitive architecture for joint action, which highlights mechanisms of prediction, monitoring, and “coordination smoothing” (Vesper et al. 2010).³ The two traditions often feel rather loosely connected: in principle, accounts of joint intention and accounts of subpersonal cognitive mechanisms should be mutually constraining and mutually informing; but in practice they have been pursued independently, and it is not always clear how they constrain and inform each other. There is a need in the joint action literature for a bridging theory: a theory of how the agent’s subpersonal mechanisms constrain and guide the joint intentional actions in which the agent can participate.

In the case of individual action, the gap between theories of intention and of subpersonal action control can be bridged by a theory of *know-how*. While there is currently little agreement as to the nature of know-how, all sides can agree that know-how plays a useful bridging role in a theory of action: know-how is a personal-level state which arises from subpersonal mechanisms of skilful action control, and which constrains and guides the agent’s intentional actions. In much the

¹ Quoted in Stix (2014, p. 77).

² For alternative theories of joint intention, see Gilbert (1989, 2000), Tuomela and Miller (1988), Tuomela (1995, 2007) and Schweikard and Schmid (2013).

³ See also Sebanz et al. (2006), Sebanz and Knoblich (2009) and Knoblich et al. (2011).

same way, the gap between theories of joint intention and of subpersonal mechanisms of joint action control can be bridged by a theory of *joint know-how*. Just as individual intentional action exercises individual know-how, joint action typically exercises joint know-how.⁴ When two people intentionally carry a log together, they exercise joint knowledge of how to carry to a log together. When two people dance the quickstep, without stepping on each other's toes or otherwise appearing badly coordinated, they exercise joint knowledge of how to dance the quickstep. When two people row a coxless pair, moving smoothly through the water without veering from side to side, they exercise joint knowledge of how to row together (my focal example in this paper, illustrated in Fig. 1 by Thomas Eakins's painting, "The Biglin Brothers Racing"). In each case, the agents' joint know-how depends on their subpersonal mechanisms and guides their joint intentional actions, suggesting an important place for a theory of joint know-how within an overarching theory of joint action.

For different reasons, epistemologists should also want a theory of joint know-how. One of the most visible debates in recent epistemology concerns the reducibility, or otherwise, of individual know-how to individual propositional knowledge.⁵ The view that know-how is reducible in this way is known as intellectualism about know-how. One can pose an analogous question in the joint case: *does joint know-how reduce to joint knowledge of a proposition?* If we can construct adequate accounts of joint know-how and joint propositional knowledge on which the former turns out to reduce to the latter, the intellectualist will be able to offer a unified account of know-how applicable to both the individual case and the joint case. That will be a mark in its favour. By contrast, if the intellectualist cannot plausibly account for joint know-how as a species of joint propositional knowledge, that is a mark against intellectualism in the individual case too, since it seems reasonable to assume (as a working hypothesis) that joint and individual know-how are not fundamentally different kinds of state, one with propositional content and the other without.

My goal in this paper is to begin the groundwork for a theory of joint know-how by examining the relationship between joint and individual know-how in paradigm cases of joint action. My question, in short, is the following: *in such cases, what is the relationship between the know-how exercised jointly by the two cooperating agents and the know-how each agent individually exercises?* For example, what does it take, in regard to the agents' individual know-how, for two agents to know how to row a coxless pair together? I will set out some basic constraints on an acceptable answer to this question (Sect. 2), argue that a simple account does not meet the constraints (Sect. 3), and construct and defend a more complex account (the "active mutual enablement" or AME account) that does (Sects. 4–6). The thought behind the AME account is that, just as joint intention (for Bratman)

⁴ I leave open the possibility that some cases of joint action do not exercise joint know-how. This may be true of cases of "prepackaged cooperation" in Bratman's (1992) sense. This possibility is discussed at the end of Sect. 2 in the main text.

⁵ For recent introductions to the debate, see Fridland (2015) and Pavese (2016). See the papers collected in Bengson and Moffett (2011) for a sense of the main positions.



Fig. 1 Thomas Eakins, “The Biglin Brothers Racing”, 1872. John Biglin sits in the stern seat (on the left of the painting); Barney Biglin sits in the bow seat (on the right of the painting). Public domain reproduction courtesy of the National Gallery of Art, Washington, DC

consists in a complex structure of individual intentions with interlocking objects, joint know-how likewise consists in a complex structure of individual know-how states with interlocking objects.

To be clear, a theory of the relation between joint know-how and individual know-how is not a *complete* theory of joint know-how, because it does not provide a theory of the intrinsic nature of know-how. It is best viewed as providing a detailed set of constraints on any unified theory of know-how that aims to encompass both the individual and joint cases. Any such theory, whatever it says about the intrinsic nature of know-how, must be such as to preserve the relation between individual and joint know-how captured by the AME account. I will argue, however, that this presents important obstacles to the development of an intellectualist theory of joint know-how (Sect. 7). One account of joint propositional knowledge (a summativist account, on which a pair of agents knows only what both agents individually know) is incompatible with any such reduction (Sect. 8). A different account of joint propositional knowledge (a distributive account, on which two agents can jointly know propositions that neither individually knows) is more conducive to the reduction of joint know-how to joint propositional knowledge (Sect. 9), but significant problems for such a reduction remain, and they push an aspiring intellectualist about joint know-how to embrace a substantial decoupling of joint knowledge from individual belief (Sect. 10).

My focus here will be on simple, two-agent cases, such as rowing a coxless pair. I will not consider larger “group agents” (such as corporations and institutions) of the sort discussed by List and Pettit (2011), nor will I consider cases of cooperation among large numbers of individual agents. I start with simple cases because my aim is to undertake the groundwork for a theory of joint know-how, and I take it to be in

two-agent cases that the relation between individual and joint know-how is easiest to apprehend (cf. Bratman 2014 on joint intention). Nevertheless, my hope is that the account constructed here will scale-up to at least some cases of larger scale cooperation. I will comment very briefly on this issue in the conclusion (Sect. 11).

2 Three constraints

Before building a positive account of the relation between joint and individual know-how, it will help to put in place a simple framework and some basic constraints. Suppose two agents, S_1 and S_2 , undertake a joint action J . Suppose the structure of J is such that each agent has a well-defined *role*, consisting of an individual action or series of actions that they must perform individually if J is to be performed jointly. This may not be true of all joint actions, but I take it to be true of paradigm cases. Consider, for example, the Biglin brothers rowing a coxless pair together: the role of John Biglin is to row in the stern seat, and the role of Barney Biglin is to row in the bow seat. The assignment of roles to agents is fixed, once they have taken their seats. Let us denote the roles adopted by S_1 and S_2 , respectively, as j_1 and j_2 .

The first constraint, Distribution, is the following:

Distribution: If S_1 and S_2 jointly know how to do J , it need not be true, of either agent, that he individually knows how to perform both roles in J .⁶

The example of the Biglin brothers intuitively motivates this constraint. For two agents to know how to row a coxless pair together, it need not be the case that both agents know how to row in both seats. It may be that one knows how to row in the bow seat, the other knows how to row in the stern seat, and neither knows how to row in the favoured position of the other.

The second constraint, Tether, demands that joint know-how is nevertheless constrained by the agents' individual know-how. For example, if we have two agents, neither of which knows how to row in either seat of a coxless pair, regardless of the circumstances, then it cannot be true of this pair of agents that they know how to row a coxless pair together. That is:

Tether: If neither S_1 nor S_2 knows how to perform any of the roles in J under any conditions, then S_1 and S_2 do not jointly know how to do J .

To be clear, Tether should not be read as requiring, of either agent, that he knows how to perform his role *in the absence of the other agent*. This would be too strong a requirement. Consider, for example, two agents lifting a log: they may know how to lift the log together by lifting one side each, even though, owing to the weight of the log, neither agent knows how to lift his side in the absence of the other. What Tether should be read as requiring is that at least one agent knows how to do at least one of

⁶ Throughout this article, I use male pronouns for the sake of continuity with my central example of the Biglin brothers.

the roles in J in at least one set of conditions, where these conditions can make reference to the other agent's actions. For example, if one agent knows how to lift his side of the log *provided the other agent simultaneously lifts his side*, this is enough to satisfy Tether as I intend it to be read.

The third constraint is intended to capture the explanatory role of joint know-how. I take it to be a central feature of know-how (whether individual or joint) that it helps explain an agent's success in executing his intentions. Admittedly, not all cases of successful execution are explained by the agent's know-how. Sometimes, in the individual case, an agent's success is explained by aid from others: for example, suppose, by cutting all the right wires, I successfully execute an intention to defuse a bomb, but only because a bomb defusal expert provided detailed instructions by phone.⁷ This is a case of success in executing my intention, but it does not manifest know-how on my part beyond that involved in cutting a wire. In other cases, an agent's success is accidental: for example, suppose I again defuse a bomb, this time without help. My success on this occasion is a matter of pure luck: I just happened to cut all the right wires. Although know-how may sometimes be explanatory of one-off successes, it is not in this case.

We should not, then, say that know-how *always* explains an agent's success in executing his intentions. What we should say instead is that when an agent *reliably* succeeds in performing some intentional action, and does so without aid from others, the agent's know-how helps explain this pattern of reliable execution. Accordingly, I suggest that the explanatory role of joint know-how is to provide an explanation of why two agents reliably succeed, without external aid, in executing a joint intention to do J . I will call this constraint Explanation:

Explanation: If S_1 and S_2 jointly know how to do J , then this provides an explanation of why they reliably succeed, without aid from others, in executing their joint intention to do J .⁸

It is worth emphasizing that this conditional is not a biconditional. If two agents jointly know how to perform some activity, this should be explanatory of their reliable success in performing that activity; but there is no implication that, if two agents reliably succeed in executing a joint intention without aid from others, then joint know-how must feature in an explanation of their success.

I therefore leave open the possibility that some cases of reliably successful joint action do not manifest joint know-how. For example, if two agents execute a joint intention to paint a house, but do so by painting separate rooms on separate days, their success can be sufficiently explained by citing prior planning, plus the fact that they individually know how to paint a room. In the terminology of Michael Bratman

⁷ This example is based on a similar example in Setiya (2012).

⁸ In contrast to Distribution and Tether, Explanation takes for granted the notion of a joint intention. This leads to the concern that it may be hard to assess whether a putative account of joint know-how satisfies this constraint unless we first commit to a particular account of joint intention. In practice, however, we will find that we do not need to commit to any particular account of joint intention to put the Explanation constraint to work. We need only assume that a joint intention is present in paradigm cases of joint action, such as rowing together.

(1992, p. 339), this is a case of “prepackaged cooperation” rather than “shared cooperative activity”. My focal cases in this paper are cases of shared cooperative activity, not prepackaged cooperation. I take it that shared cooperative activity usually (perhaps always) exercises joint know-how, whereas prepackaged cooperation does not. For Bratman, the difference between shared cooperative activity and prepackaged cooperation is marked by the presence of “mutual responsiveness in action” in the former but not the latter. Mutual responsiveness in action is related to joint know-how, but the relation between them is subtle; I will revisit this issue in Sect. 5.

3 A simple account (and its inadequacy)

As a starting point, let us consider a simple proposal:⁹

Simple account:

S_1 and S_2 jointly know how to do J if and only if:

- (1) S_1 knows how to perform his assigned role, j_1 , in conditions c_1 ;
- (2) S_2 knows how to perform his assigned role, j_2 , in conditions c_2 ;
- (3) c_1 and c_2 are mutually compatible.

The simple account satisfies Distribution and Tether. As previewed in the discussion of Tether, the conditions c_1 and c_2 may make reference to the actions of the other agent. For example, it may be that John knows how to row in the stern seat only when a competent rower is sat in the bow seat, and vice versa. The agents need not know how to do their parts in isolation, or with incompetent partners. What is crucial is that the conditions each requires for the exercise of his knowledge of how to do his part are compatible with each other.¹⁰

Conditions (1)–(3) in the simple account seem to capture an important aspect of what it is for S_1 and S_2 to know jointly how to do J , but they do not suffice. Consider again the example of the Biglin brothers. Suppose S_1 is the rower in the stern seat (John) and S_2 is the rower in the bow seat (Barney). The two roles (j_1 and j_2) are rowing in the stern seat and rowing in the bow seat. Suppose that each agent knows how to perform his part under mutually compatible conditions, so that conditions (1)–(3) are satisfied. Does this suffice for S_1 and S_2 's knowing how to row together?

It does not. Intuitively this is because, even granting that (1)–(3) obtain, it may be that as soon as S_1 and S_2 begin to row, the boat starts lurching and veering

⁹ An even simpler account would omit c_1 , c_2 , and condition (3). But such an account would lead us to attribute joint know-how even in cases in which the agents know how to perform their roles only in mutually incompatible conditions and would be clearly inadequate for this reason.

¹⁰ In some cases, an agent may require instruction (or some other form of aid) from the other. The simple account allows for this possibility (as does the AME account, below). It is simply that, in such cases, the specification of c_1 and/or c_2 should incorporate the required aid. This also makes room for the possibility that two agents may jointly know how to do J even though neither knows how to do his part without aid from the other.

waywardly, owing to their performing their roles at different rates or in otherwise uncoordinated ways, and this seems in conflict with the idea that they know how to row together. The conditions have missed the crucial point that the agents must know how to coordinate their individual performances.

We can give a more principled reason why these conditions do not suffice by referring back to Explanation. If S_1 and S_2 's joint know-how consisted merely in the conjunction of S_1 's knowing how to row in the stern seat and S_2 's knowing how to row in the bow seat under mutually compatible conditions, their joint know-how would offer no explanation for their reliably successful execution of their joint intention to row together. This is because reliable success here, and in shared cooperative activity more generally, depends on mutual coordination, yet conditions (1)–(3) offer no explanation of mutual coordination. So these conditions do not, by themselves, yield an account that satisfies Explanation.

In short, the simple account satisfies Distribution and Tether but not Explanation. How can we amend the account so that it does? It may be tempting to simply add to conditions (1)–(3) a fourth condition:

- (4) S_1 and S_2 know how to coordinate their performances.

This seems to capture the crucial missing ingredient. However, it accounts for the agents' joint knowledge of how to do J in terms of a further, unanalysed piece of joint know-how. Given our goal of elucidating the relationship between joint and individual know-how, this is unsatisfying. The challenge we face is that of explicating condition (4) in terms of further facts about what the agents individually know how to do.

4 Knowing how to coordinate

What is involved in two agents knowing how to coordinate in a joint action? First, a remark on how I think we should *not* approach this question. There is a substantial game-theoretic literature on coordination, in which we find, for example, the notion of a coordination equilibrium. A coordination equilibrium is, roughly, a special type of Nash equilibrium such that no agent stands to gain an increased payoff when any agent unilaterally switches strategy. Might the sort of coordination that characterizes shared cooperative activity be described as the agents attaining a coordination equilibrium?

I suggest not, for two reasons. First, the notion of a coordination equilibrium is too thin for my purposes: two agents may achieve such an equilibrium in a game in which their strategy profile bears little resemblance to a joint action. For example, suppose two agents, both driving their cars on the left, pass each other in the night while travelling in opposite directions, without evening so much as a glance in each other's direction. This might be modelled as a coordination equilibrium, but it is not displaying the type of coordination characteristic of joint action. Second, there is also a sense in which the notion is too demanding, since it requires that coordination leads to both agents maximizing their personal payoff. Yet personal payoffs need

not be maximized in a joint action: indeed, there may be no meaningful payoff at stake. Along these lines, Brian Skyrms (2010, p. 57) writes of the “game of inconsequential actions”, in which all payoffs are zero, and writes “I believe that much of life has this structure”. The type of coordination that concerns me here is compatible with the joint action being inconsequential in this sense.¹¹

How, then, should we approach this question? As we noted in Sect. 2, Bratman (1992, 2014) describes “mutual responsiveness in action” as the mark of shared cooperative activity, and this is a starting point: coordination requires that each agent in some sense responds to the other. However, as Bratman’s critics have noted, he does not offer a detailed explication of mutual responsiveness (Gold and Sugden 2007).¹² To put some flesh on the bones of this idea, let us first note three obvious features of the sort of mutual responsiveness that leads to successful joint action. First, the responses in question are of a particular kind: they take the form of *adjustments* to one’s own performance in light of what one’s cooperative partner is doing. Second, not just any adjustment will do: cooperating agents adjust their behaviour so as to *avert the failure* of the joint action. For example, if two rowers intend to row in a straight line but are rowing waywardly, they exercise their know-how when one adjusts his stroke rate to restore synchronicity, removing the waywardness. Third, these adjustments are not the result of prior planning, but rather come as a result of *predicting* and *monitoring* what the other agent is doing in real time.¹³

As the psychological literature on joint action shows, the agents’ monitoring of each other may take various forms (Vesper et al. 2010). It may take the form of direct perceptual monitoring of the other agent’s behaviour, but it may also be mediated by cues and signals that indicate aspects of the other agent’s behaviour. Indeed, if one compares the performance of experts at a joint action to the performance of two novices, it is a common phenomenon that the novices make much heavier use of direct perceptual monitoring than the experts. For example, two novice dancers may need to watch each other’s feet, and two novice pianists may need to watch each other’s hands as they play a duet. Experts often do without this extreme form of direct monitoring, instead relying more heavily on signals and cues.

So far, we have the idea that coordination in the context of joint action consists in the failure-averting adjustment of the agents’ performances as a result of their

¹¹ The literature on “team reasoning” offers an alternative to orthodox game theory that aims to capture the distinctive sort of practical reasoning that occurs in the context of joint action (e.g. Bacharach 2006; Gold and Sugden 2007). Nothing I say here is in tension with that approach, because I have a different aim: my aim is not to characterize the practical reasoning that leads rational agents to decide to coordinate, but rather to characterize the know-how exercised when they coordinate.

¹² To illustrate mutual responsiveness in action with reference to his focal example of a duet, Bratman (1992, p. 339) writes: “I listen closely to when and how you come in, and this helps guide my own singing; and you are similarly responsive.” This foreshadows the “monitoring” component of the account I develop here, and perhaps the “adjustment” component, but not the “active enablement” component.

¹³ “The Biglin Brothers Racing” is an illustration of this point. On close inspection, we see the brothers are slightly out of sync: Barney is slightly ahead of John in his stroke. In the painting, we see him watching John’s oar—indirectly monitoring his behaviour—with a view to adjusting his own movements to restore synchronicity.

prediction and (direct or signal-mediated) monitoring of each other's behaviour. However, something important is missing: the effort invested by an agent to make himself *such as to be easily predicted and monitored*. Stephen Levinson once remarked that “there is an extraordinary shift in our thinking when we start to act intending that our actions should be coordinated with—then we have to design our actions so that they are self-evidently perspicuous” (1995, p. 241).¹⁴ I take this to be an important insight, and one our account has so far missed. Coordination is not simply a matter of each agent making adjustments to their performance in response to that of the other. It is also a matter of each agent, from the outset, acting in such a way as to *actively enable* the other agent to make such adjustments. In other words, each agent performs his role in a way that is *actively coordination-enabling* for the other.

What is involved in an actively coordination-enabling performance? Here, Vesper and et al. (2010) discussion of “coordination smoothing” is particularly helpful. Vesper et al. note that in many cases, including crew rowing, prediction and monitoring is substantially easier if the target is able to make his own performance less variable by, for example, rowing at a steadier rate. Feigean et al. (2017) show that, when two skilled individual rowers learn to row a coxless pair for the first time, the rower in the stern position learns to become less variable, producing slower, larger-amplitude sweeps of the oar at a more constant rate. In other cases, what matters is the making salient of certain aspects of the performance through subtle signals. Vesper et al. (2010, p. 1001) give the example of two pianists playing a duet, using synchronized head movements and pronounced finger raising as a way of keeping time with each other when they are unable to hear the piano (Goebel and Palmer 2009).

Inevitably, the means by which predictability and perspicuity is achieved in a joint action will depend a great deal on the nature of the action in question. I suspect all that can be said in general about actively coordination-enabling performance is that it consists of an agent making predictable and perspicuous to the other agent those aspects of his performance that may call for responsive adjustment—and to do so not by accident but intentionally, in a way that is appropriately sensitive to the other agent's ability to predict, monitor and respond. For our purposes, the following working definition will be enough:

S_1 's performance of j_1 is *actively coordination-enabling* for S_2 just in case the features of the performance that may call for responsive adjustment from S_2 are intentionally made predictable and perspicuous to S_2 , either through the nature of the performance itself or through the use of signalling.

¹⁴ See also Tomasello (2014, p. 58), who also quotes this sentence and further expands on Levinson's insight.

5 The active mutual enablement (AME) account

These considerations lead to a proposal I will call the *active mutual enablement* (AME) account of joint know-how:

AME:

S_1 and S_2 jointly know how to do J if and only if:

- (1a) S_1 knows how (in conditions c_1) to perform j_1 in a way that is actively coordination-enabling for S_2 .
- (1b) S_1 knows how to predict, monitor and make failure-averting adjustments in response to S_2 's performance of j_2 , provided S_2 performs j_2 in a way that is actively coordination-enabling for S_1 .
- (2a) S_2 knows how (in conditions c_2) to perform j_2 in a way that is actively coordination-enabling for S_1 .
- (2b) S_2 knows how to predict, monitor and make failure-averting adjustments in response to S_1 's performance of j_1 , provided S_1 performs j_1 in a way that is actively coordination-enabling for S_2 .
- (3) c_1 and c_2 are mutually compatible.

The AME account claims that joint know-how is reducible to an interlocking package of individual know-how states with distinctive and complex objects: each agent knows how to perform his role in an actively coordination-enabling way for the other agent while predicting, monitoring, and making responsive adjustments in response to the other agent's behaviour.

I contend that the proposed account satisfies our three basic constraints. Like the simple account, the AME account satisfies Distribution and Tether, and I submit that it also satisfies Explanation. When agents reliably execute a joint intention to perform a joint action, it is the relational structure between their individual know-how states described by AME account—each agent knowing how to actively enable responsive adjustments from, while making responsive adjustments to, the performance of the other agent—that explains the reliability of their success in the face of unforeseen obstacles and challenges.

The complementarity between the AME account and Bratman's (1992, 1993, 2014) theory of shared intention should now be clear. Bratman accounts for shared intention in terms of individual intentions with distinctive, interlocking objects involving relations to the other agent's intentions. Very roughly, each intends that we do J in accordance with and because of what Bratman terms meshing subplans. Similarly, the AME account characterizes joint know-how in terms of individual know-how states with distinctive, interlocking objects involving relations to the other agent's know-how. Although I have not assumed the correctness of Bratman's account of joint intention in constructing the AME account, the two theories are intended to be complementary.

One disanalogy with Bratman's theory is worth noting: Bratman's conditions for shared intention include a common knowledge requirement, but the AME account contains no such requirement. It need not be common knowledge that the conditions

for joint know-how are satisfied. In principle, two agents could jointly know how to do J even though neither knows that the other satisfies the relevant conditions. There is room for debate about whether a common knowledge requirement is justified even in the case of shared intention.¹⁵ However, even if we grant that it is, there is no justification for a corresponding requirement regarding joint know-how. If John secretly harbours the false belief that Barney is incompetent, this does nothing to change the fact that, in reality, they know how to row together.

We can also revisit, as promised at the end of Sect. 2, the relationship between joint know-how and Bratman's notion of "mutual responsiveness in action". Joint know-how will often be *manifested* in the form of mutual responsiveness: agents will exercise their joint know-how in part by responding to each other. It is possible, however, for two agents to know how to do something without their know-how being manifested in this way. Sometimes everything goes so smoothly that no adjustments are required. If a performance is routine and the conditions favourable, monitoring may not even be required. Joint know-how, like individual know-how, can be present but latent. More subtly, agents may manifest their joint know-how in a non-standard way by intentionally performing a joint action badly (e.g. for instructional purposes), demonstrating the sort of non-accidental miscoordination that only two agents who know how to coordinate could reliably achieve.¹⁶

Moreover, mutual responsiveness need not be a manifestation of joint know-how, because it may arise even in cases in which the agents do *not* know how to perform the joint action they are attempting. For example, two singers who *don't* know how to sing a duet together may nonetheless listen carefully to each other and adjust their performances—but not in a way that reliably averts failure. Joint know-how and mutual responsiveness in action are related but distinct.

6 Objections: narrowness, symmetry and regress

One objection to the AME account is that its conditions are not satisfied in some cases we might intuitively describe as involving joint know-how. These are cases of prepackaged cooperation (in Bratman's sense) in which the agents do not predict, monitor or respond to each other, but nevertheless cooperate in a sufficiently intricate way that an ascription of know-how to the pair or group seems intuitive. Consider, for example, two singers singing a duet together. Suppose that each knows how to perform her part only if the other performs her part perfectly (so that no adjustments to the tempo or the pitch of the other singer are necessary) and that each singer knows how to perform her part perfectly.¹⁷ Even though they lack joint

¹⁵ Blomberg (2016) argues that the requirement was never well motivated, but see Paternotte (2013) for a defence of the requirement.

¹⁶ As Stanley and Williamson (2017) observe, it is a general feature of skills that they can be manifested in non-standard ways, such as intentionally poor performance.

¹⁷ This example, suggested by an anonymous reviewer, is based on Bratman's (1992, pp. 336–337) case of the "unhelpful singers".

know-how in the sense of the AME account, is it not intuitive to say they know how to sing the duet together?

It is initially intuitive, I agree. However, small modifications to the example make the ascription of joint know-how less intuitive. Suppose the two singers are in adjoining recording studios, unable to hear or see each other. If the intuitiveness of ascribing joint know-how does not diminish, suppose that they record their parts on separate days, and in separate locations. The modified example seems like a clear case of two individuals independently exercising their individual know-how. Spatiotemporal proximity, mutual audibility or mutual visibility are not necessary for prepackaged cooperation, even prepackaged cooperation of an intricate and interweaving nature, but they seem to make a difference to the intuitiveness of an ascription of joint know-how. Why is that? Because, I suggest, despite the initial plausibility of joint know-how ascriptions in some cases of prepackaged cooperation, joint know-how requires knowledge of how to enable, monitor and adjust to the other agent. I suspect that the ascription of joint know-how initially seems plausible because we conflate a case in which the singers genuinely do not know how to enable, monitor, or adjust to each other with a superficially similar case in which this know how is present but latent, because the performances are perfectly executed.

A different objection derives from the fact that the AME account, like Bratman's theory of joint intention, makes *symmetric* demands of the two agents' individual know-how. That is, although the nature of the actions j_1 and j_2 may differ considerably, the *form* of the requirements on each agent's know-how is the same: each must know how to do his part in an actively coordination-enabling way while predicting, monitoring and adjusting to the other agent.¹⁸ This leads to a potential concern: is this symmetry always appropriate?

To motivate the concern, consider again the Biglin brothers. Suppose that John, who sets the stroke rate, knows how to do his part in an actively coordination-enabling way (he rows steadily, clearly signalling the rate at which he intends to row) but does not know how to make adjustments to Barney. Suppose, meanwhile, that Barney knows how to adjust to John—he reliably matches the rate John sets—but does not know how to actively enable adjustments from John. Would this asymmetric distribution of know-how suffice for them to know jointly how to row together? If so, the AME account requires too much: the “one-way coordination” present in this example is already enough.

I contend, however, that one-way coordination is not enough. What is true about the case of the Biglin brothers is that, because both are experts, there is no need for John to monitor Barney's performance directly. Nonetheless, they would not reliably succeed at executing this demanding joint skill if John were wholly unresponsive to Barney. Barney must know how to communicate any change in circumstances (e.g. communicating the need for a change in the stroke rate), and John must know how to adjust his performance in light of the information received.

¹⁸ To put this more precisely, the demands are symmetric in the sense that one could permute the agents (swapping S_1 for S_2), the individual actions (swapping j_1 for j_2) and the conditions (swapping c_1 for c_2) in the AME account and thereby arrive at a notationally equivalent account.

If Barney were wholly uncommunicative, or if John were wholly unable to respond to information from Barney, I maintain that they would not possess joint know-how. Their success in executing their joint intention would be reliant on fortuitous circumstances that did not expose the fragile nature of their coordination.

A third objection concerns an apparent threat of regress that arises when we ask how joint know-how is acquired. For how can two agents learn how to do *J* together unless they already know how to perform some other joint action together? For example, two capable rowers will normally require a period of joint practice to learn how to row a coxless pair together.¹⁹ If they know immediately how to row together, this can only be because they have a long history of joint practice with other, similar partners. It is a virtue of the AME account, not a defect, that it explains why joint know-how tends to be cultivated through joint practice. The agents may know how to perform their individual roles from prior individual practice, but this is not enough to meet the requirements of the account: each must learn how to make responsive adjustments to the other, and each must learn how to perform his role in a way that is actively coordination-enabling for the other. While this mutual attunement of each agent's individual know-how to that of the other might in principle be achieved without joint practice, especially when the agents are already attuned to other partners who behave in similar ways, it is no surprise that joint practice is often needed.

But here is where the puzzle arises: joint practising of *J* is itself a joint action. Given this, how can two agents who do not already know how to do *something else* together—namely, how to practise doing *J*—ever come to possess the joint knowledge of how to do *J*? How, in turn, can they ever come to possess the joint knowledge of how to practise doing *J* without first knowing how to do something else together—roughly, how to practise doing some of the *elements* of *J*? If the pattern continues indefinitely, it becomes impossible to see how two such agents could ever make the initial transition from possessing no joint know-how to possessing some.

It must be, therefore, that the pattern does not continue indefinitely: instead, it bottoms out in basic forms of joint know-how that are *not* acquired through joint practice, but instead arise from innate or individually learned aspects of general practical competence.²⁰ This seems less mysterious once one observes that knowing how to practise doing *J* is less of an achievement than knowing how to do *J*. The success conditions for practising *J* are more relaxed than the success conditions for doing *J*, which makes it easier for the agents' individual performances to suffice for success, and also makes it easier for the agents' adjustments to qualify as failure-averting. For practising doing some element of *J*, the success conditions are more relaxed still, and so on. Eventually, we arrive at basic cooperative activities with sufficiently relaxed success conditions that any generally competent agents who form a joint intention to do such an activity can immediately be said to know how to

¹⁹ For a detailed account of how practising in a coxless pair modifies the technique of the rowers, as they learn to be actively coordination-enabling for each other, see Feigean et al. (2017).

²⁰ For discussion of the notion of "general practical competence" on which I rely here, see Railton (2009).

do it, without any need for prior practice. For example, two novice rowers who do not even know how to practise rowing together might still, by virtue of their general practical competence, know how to practise raising and lowering their oars in time with one another.

The basic platform of joint know-how that arises from the agents' general practical competence thus allows for the acquisition, through joint practice, of more demanding know-how states. This occurs by means of a familiar ratcheting, whereby mastering a given level of performance in some joint action makes it possible to acquire joint knowledge of how to perform at a slightly higher level, and so on. Two novice rowers who have mastered the synchronized raising and lowering of their oars can move on to practising whole strokes, while two rowers who know how to perform whole strokes in time with one another can start learning how to coordinate their movements more precisely. It is a virtue of the AME account that it allows for this kind of ratcheting.

7 Intellectualism about joint know-how: the core commitment

In constructing the AME account, I made no assumptions about whether know-how should, or should not, be interpreted as a species of propositional knowledge. But having arrived at this account, we can now ask: can joint know-how be interpreted as joint knowledge of a proposition? The picture here is complicated by the absence of any agreed theory of joint propositional knowledge. We can ask, however, which views about joint propositional knowledge are compatible with the reduction of joint know-how to joint knowledge of a proposition, and which are not. I will consider (in Sects. 8–10) two such views, both of which share the assumption that joint propositional knowledge depends closely on individual propositional knowledge, but which disagree about the nature of the dependence.

In the case of individual know-how, the standard bearer for intellectualism is Jason Stanley (2011), who proposes that an agent knows how to do an individual action if and only if the agent knows, of some way w with which the agent is acquainted under a practical mode of presentation, that w is a way in which one could successfully do that action. To exercise one's know-how is to implement a way one knows could lead to success. In slogan form, know-how is knowledge of counterfactual success.²¹

We can construct an analogue of Stanley's intellectualism in the case of joint know-how. Let (w, v) represent a joint way of performing J , comprising two individual ways, denoted by w and v , of performing the two individual roles in J . We should allow here that w and v incorporate ways of making one's performance coordination-enabling for the other agent, as well as ways of predicting, monitoring and adjusting to the other agent's behaviour. In other words, w and v are to be understood not as ways of performing individual actions without regard to the other

²¹ For recent discussion of this and related proposals, see Setiya (2012), Fridland (2013, 2014), Carter and Pritchard (2015), Santorio (2016) and Pavese (2015, 2017).

agent, but as ways of fully satisfying requirements (1) and (2), respectively, of the AME account. The core commitment of Stanley-style intellectualism as applied to joint know-how is that joint know-how is joint knowledge of counterfactual success. We can express this core commitment as follows:

Intellectualism about joint know-how: S_1 and S_2 know jointly how to do J if and only if S_1 and S_2 jointly know, of some joint way (w, v) , that (w, v) is a way of successfully doing J .

I have omitted a joint-level analogue of Stanley's requirement that the agent is acquainted with the way under a "practical mode of presentation" (or "guise", or "way of thinking"). This is because, even if this notion can be made sufficiently precise in the case of individual action, it remains unclear what the joint-level analogue of this requirement should be, and this complication does not matter to my arguments in Sect. 8.²² I will, however, revisit this omission in due course, since it provides the intellectualist with a possible escape route from the objection presented in Sect. 9.

8 Summative intellectualism

The core commitment of intellectualism about joint know-how can be combined with different views on the nature of joint propositional knowledge to yield substantially different versions of intellectualism. One version, which I will here call *summative intellectualism*, conjoins the intellectualist's core commitment with a summativist view of joint propositional knowledge, on which a pair of agents can be said to know only that which both agents know individually. The result is the following view:

Summative intellectualism: S_1 and S_2 jointly know how to do J only if both agents *individually* know, of some joint way (w, v) , that (w, v) is a way of successfully doing J .

There is an intuitive tension between summative intellectualism and Distribution, which I introduced in Sect. 2 as a basic constraint on any acceptable account of joint know-how. Recall that, according to Distribution, joint know-how does not require that each agent knows how to perform both roles in the joint action. Assuming for the sake of argument that intellectualism about individual know-how is correct, Distribution implies that, for two agents to know jointly how to do J , it need not be the case that either agent knows, for both roles, a way of performing that role. This is not logically incompatible with summative intellectualism, but an incompatibility is generated by adding one further plausible assumption that I will call *Decomposability*: an agent individually knows that (w, v) is a way of successfully

²² For criticism of the notion of a "practical mode of presentation", see Noë (2005), Stalnaker (2012) and Glick (2015). For a recent defence, see Pavese (2015).

doing J only if the agent individually knows that w is a way of successfully performing j_1 and that v is a way of successfully performing j_2 .

In short, the argument against summative intellectualism, from Distribution, Decomposability and individual-level intellectualism, is as follows:

- (1) *Intellectualism about individual know-how*: S_1 knows how to do j_1 if and only if S_1 knows, of some way w , that w is a way of successfully performing j_1 .
- (2) *Decomposability*: S_1 knows that (w, v) is a way of successfully doing J only if S_1 knows that w is a way of successfully performing j_1 and that v is a way of successfully performing j_2 .
- (3) *Distribution*: If S_1 and S_2 jointly know how to do J , it need not be true, of either agent, that he individually knows how to perform both roles in J .
- (4) Therefore (from 1 to 3): If S_1 and S_2 jointly know how to do J , it need not be true, of either agent, that he individually knows that (w, v) is a way of successfully doing J .

This argument puts the summative intellectualist in a difficult position. To be of any plausibility, summative intellectualism about joint know-how must be combined with intellectualism about individual know-how. Thus the summative intellectualist must choose between two unattractive options: reject Distribution, a constraint motivated by reflection on paradigm cases of joint action (see Sect. 2); or else reject Decomposability, and maintain that an individual can know a proposition concerning the counterfactual success of a joint way of acting without knowing the corresponding propositions about the counterfactual success of the individual ways of acting from which the joint way is composed.

Note here that this problem in no way depends on the precise details of the AME account. The source of the problem is Distribution, and it will arise given any account of the relationship between joint and individual know-how that satisfies Distribution.

9 Distributive intellectualism

I take it, then, that a more promising option for the intellectualist is to reject the summativist view of joint propositional knowledge in favour of a view on which joint propositional knowledge, like joint know-how, can be distributed between the participants in a joint action. The intuitive idea here is that, for two agents to jointly know that p , it need not be the case that either one individually knows that p : it may be enough that a hypothetical observer with full access to the propositional knowledge of both individuals could reasonably infer p from the set of truths known to at least one of them. This is not to deny that joint propositional knowledge is closely tied to individual propositional knowledge, but it allows the ties to be somewhat looser than the summativist view allows. This leads to a view which I will call *distributive intellectualism*:

Distributive intellectualism: S_1 and S_2 jointly know how to do J if and only if the proposition that (w, v) is a way of successfully doing J can be reasonably inferred from the set of propositions individually known by S_1 or S_2 .²³

That joint propositional knowledge can indeed be distributed between agents in this way is made independently plausible by cases of joint memory (Barnier et al. 2008; Sutton et al. 2010).²⁴ For example, Harris et al. (2014) investigated, through tests and interviews, the extent to which remembering as a couple facilitates or inhibits recall. Joint remembering, while not always beneficial, appears to facilitate the recall of episodic and qualitative detail. For example, in one interview, a couple were able to remember the name of a show they had seen on a particular occasion, even though neither could remember the name individually (see Harris et al. 2014, pp. 291–292). The husband remembered that the show they had seen starred John Hanson, and the wife remembered that they had seen a show starring John Hanson that was called *The Desert Song*. We can plausibly interpret this as a case in which the couple knew jointly that they had seen *The Desert Song* on the occasion in question, but neither knew individually until they had actively integrated their individual knowledge. Examples of this sort motivate (independently of any considerations relating to know-how) an account of joint propositional knowledge that does not require a jointly known proposition to be individually known by either agent, and that allows a proposition to be jointly known when it may be inferred from the propositions each agent individually knows.²⁵

Distributive intellectualism is thus more credible at face value than summative intellectualism. Agents who satisfy all the conditions of the AME account will, if we assume the truth of individual-level intellectualism, individually know a package of propositions concerning ways of performing their roles, ways of enabling coordination in each other, and ways of monitoring and adjusting to each other's performance. It is not implausible that, when all these conditions are satisfied, the proposition that (w, v) is a way of successfully doing J could be inferred from all that they individually know.

However, distributive intellectualism still faces a significant problem. The problem is that there are cases in which two agents jointly know how to perform a joint action while believing that the joint way they implement is one that could not lead to success. For a vivid example, consider *anti-phase rowing*. In traditional rowing, the rowers aim for perfect synchrony with each other. This leads to cyclical

²³ The difference between summative and distributive versions of intellectualism maps on to a more general division between summative and distributive approaches to collective knowledge and belief. For an introduction to the general issue, see Lackey (2014).

²⁴ See Bird (2014) for arguments that propositional knowledge can also be distributed across larger groups.

²⁵ An account of joint propositional knowledge that, in my view, meshes even better with these results is one on which two agents know jointly all and only those propositions which *they know how to infer*, through communication with each other, from their stock of individual propositional knowledge. However, such an account explicitly makes joint know-how prior to joint propositional knowledge. Such a view is not available to the intellectualist, who must hold that joint propositional knowledge is prior to joint know-how. Thus distributive intellectualism should not be formulated in these terms.

velocity fluctuations that undermine efficiency: the boat slows, then speeds up, then slows, and so on. A way to maintain smooth velocity is to row in perfect anti-phase. This is very difficult to achieve, but experiments using coupled rowing machines have shown that there is the potential for significant power benefits if coordination can be maintained (de Brouwer et al. 2013; Cuijpers et al. 2015; de Poel et al. 2016).

Imagine, then, the following case. Two skilled, experienced rowers row in phase together in a coxless double,²⁶ exercising their joint knowledge of how to row together. Now an experimenter asks them to start rowing anti-phase. They are highly sceptical of this instruction. They believe that, as soon as they try it, the boat will become unbalanced, wayward and slow. They sincerely believe a boat cannot be rowed like this. As a matter of fact, however, the agents fully satisfy the conditions of the AME account with respect to anti-phase rowing. By virtue of their existing skill and experience, each knows how to do his part while monitoring and adjusting to the other, and enabling such monitoring and adjustment from the other. The two rowers already know how to row anti-phase, despite disbelieving that the method can work. As a result, all goes well: they maintain anti-phase coordination, and the boat continues to glide through the water and at a smoother velocity than before.

Consider the rowers' state of mind at the moment the experimenter gives the instruction. I contend that at this moment (a) the two rowers jointly know how to row anti-phase by virtue of their existing competence and manifest this joint know-how in their subsequent action; but that (b) neither agent believes that what they are about to do is a way of rowing together that could lead to success. Indeed, both agents believe themselves, falsely, to be preparing to implement a joint way that is guaranteed to fail as a way of rowing together.

I take this example to show that a pair of skilled agents may possess joint knowledge of how to do J even though both agents individually believe, of the way in which they know how to do J , that it cannot lead to success.²⁷ The example refutes distributive intellectualism given the following additional assumption, which I will call *Exclusion*:

Exclusion: If both S_1 and S_2 individually believe that not- p , then they do not jointly know that p .

Again, the options for the intellectualist look unattractive. The distributive intellectualist may either deny Exclusion, an assumption that merely rules out an extreme kind of tension between joint and individual belief,²⁸ or else deny the possibility of two agents jointly knowing how to do something while believing that the technique they jointly implement cannot succeed.

²⁶ In a double, as opposed to a pair, both rowers hold two oars. This is more stable under anti-phase rowing than a pair.

²⁷ The possibility of such examples is related to the absence of a common knowledge requirement regarding joint know-how, as discussed in Sect. 5.

²⁸ This radical decoupling of individual from joint or collective belief has at least one notable defender: Gilbert (1987, 2005). It is not an absurd or obviously indefensible option—just an unattractive one.

10 Practical modes of presentation

The idea of a mismatch between an agent's avowed beliefs and know-how is not unique to the case of joint action. Brownstein and Michaelson (2016) consider various cases in which an agent appears to know how to perform some individual action while believing, of the way he implements, that it cannot lead to success. However, as Brownstein and Michaelson (2016, p. 2830) acknowledge, the intellectualist can seek to account for these examples by positing beliefs held under a *practical mode of presentation* that contradict the agent's avowed beliefs (Stanley and Williamson 2001; Stanley 2011, Ch. 4). I may believe of a way w , presented under a theoretical guise, that it could not lead to success, and yet know of that same way, presented under a practical guise, that it could lead to success.

Consider, for example, a way of cricket batting that involves *looking ahead of the ball rather than at it*. It is well known that cricketers almost always deny that this way of batting, presented under a theoretical guise, actually works. But eye-tracking studies show that successful batters do indeed look ahead of the ball.²⁹ The intellectualist will seek to handle such cases by arguing that the batter knows of that same way of batting, presented under a practical guise, that it is a way of batting successfully.

Let us grant for the sake of argument that this move works in the individual case. Could a similar move rescue distributive intellectualism in the joint case? We could grant that our rowers (like the cricket batters) both believe, of a joint way with which they are acquainted under a theoretical mode of presentation (i.e. anti-phase rowing), that it could not lead to success. But we could maintain that, nevertheless, they jointly know of that same way, with which they are jointly acquainted under a practical mode of presentation, that it is a way in which they could successfully row anti-phase.

To make this work, we need to say what is required for two agents to be "jointly acquainted under a practical mode of presentation" with a way of performing a joint action. If joint acquaintance requires past experience of implementing this particular joint action together, our rowers do not have it. But perhaps this is too strong a requirement. A more liberal conception of joint acquaintance requires only that each agent is individually acquainted, under a practical mode of presentation, with a way of performing his own role (where, as before, the "way" in question is understood to encompass all the actions required by the AME account, and not just performance of an action without regard for the other agent). The rowers in our example might then be said to be jointly acquainted under a practical mode of presentation with a way of successfully rowing anti-phase, despite never having done so.

This more sophisticated version of distributive intellectualism can be formulated as follows:

*Distributive intellectualism**:

²⁹ See Mann et al. (2013). This example is discussed in detail by Brownstein and Michaelson (2016, p. 2822).

- i. S_1 and S_2 jointly know how to do J if and only if they jointly know, of some way (w, v) with which they are jointly acquainted under a practical mode of presentation, that (w, v) is a way of successfully doing J .
- ii. They are *jointly acquainted* with (w, v) under a practical mode of presentation if and only if at least one agent is individually acquainted with w under a practical mode of presentation and at least one agent is individually acquainted with v under a practical mode of presentation.
- iii. They *jointly know* that (w, v) is a way of successfully doing J if and only if this proposition can be reasonably inferred from the set of propositions individually known by S_1 or S_2 .

This provides a potential escape route for the distributive intellectualist. Agents who satisfy all the conditions of the AME account may be regarded as jointly knowing, of a joint way with which they are acquainted under a practical mode of presentation, that it is a joint way that would lead to success. Strictly speaking, the principle a Stanley-style intellectualist must deny is not Exclusion as formulated in Sect. 9, but rather Exclusion*:

*Exclusion**: If both S_1 and S_2 individually believe, under a theoretical mode of presentation, that not- p , then they do not jointly know, under a practical mode of presentation, that p .

The denial of Exclusion* is not such a bitter pill for intellectualists to swallow, especially since they must already swallow a similar pill to handle cases of belief/know-how mismatch in the case of individual action. To be clear, however, the escape route is open only to the distributive form of intellectualism. The summative form remains implausible for the reasons given in Sect. 8.

11 Open questions

The AME account, and the above considerations on its relationship to intellectualism, provide a starting point for further investigation of the nature, evolutionary origin and scaling-up of joint know-how. Although I cannot do justice to these topics here, I will offer some brief reflections in the hope of generating further debate.

Let us first consider the question of origins. Can chimpanzees possess joint know-how? They undoubtedly meet some of the necessary conditions: they often know how to make adjustments to the behaviour of others in contexts such as hunting. However, I doubt that they know how to act in *actively coordination-enabling ways* in the above sense. For this requires advanced mindreading: it requires an understanding of how others are likely to respond to one's own performance—what others will see, what they are able to predict, the inferences they are likely to make, and the ways they are likely to adjust their behaviour in response. Language is not a pre-requisite, but it is one of the many ways in which humans demonstrate their vastly greater capacity for “acting so as to be coordinated with”. It seems likely that chimpanzees have some capacity for mindreading (Krupenye et al. 2016), although

this remains a source of some controversy (Heyes 2017). But it seems much less likely that they know how to act in ways that are actively coordination-enabling for their cooperative partners.

If this is right, then the evolution of joint know-how is part of the story of hominin evolution, post-dating our divergence from the other great apes. It is entangled, in ways we have not yet begun to explore, with the evolution of joint intention, mindreading and language. Unpicking these threads is an important direction for future work.

What about scaling up? I have left open the question of how the AME account may (or may not) scale up to larger feats of cooperation. The account extends naturally enough to small-scale cooperative activities involving more than two agents: for an activity involving three agents, for example, one can readily imagine an extended account requiring that each agent knows how to perform his part in a way that is actively coordination-enabling for the other two, while predicting and monitoring what the other two are doing and making failure-averting adjustments.

However, this strategy only gets us so far. In particular, the account does not straightforwardly extend to cover the know-how manifested by paradigm cases of “group agents”—institutions, corporations, and so on—of the sort analysed by List and Pettit (2011). If we say, for example, that Apple knows how to make large profits, this is not a case of know-how that may be readily explicated in terms of individual agents individually knowing how to monitor each other and make adjustments: something qualitatively different is going on here.³⁰

Nonetheless, it seems reasonable to hypothesize that there are connections between the know-how of a group agent and the joint (or shared) know-how of pairs (or larger subsets) of individuals within the group. After all, small-scale cooperative activities often form a vital component of the lower-level machinery by means of which the actions of a group agent are realized: Apple’s profit-making, for example, depends on the collaborative activities of subsets of its employees. These connections are worthy of further exploration.

Acknowledgements I am very grateful to Alexander Bird, Olle Blomberg, Lucy Campbell, Peter Dennis, Christian List, Anna Mahtani, Eliot Michaelson, Cedric Paternotte, Greg Radick, Kieran Setiya, John Sutton, an anonymous reviewer and audiences at the LSE, UCL, the University of Leeds and the University of Edinburgh for their insightful comments. This work was supported by a Philip Leverhulme Prize from the Leverhulme Trust, 2015–2018, and by a visiting fellowship at the Australian National University in 2015.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

³⁰ See Palermos and Tollefsen (2018) for a preliminary investigation of group know-how.

References

- Bacharach, M. (2006). *Beyond individual choice: Teams and frames in game theory*. Princeton: Princeton University Press.
- Barnier, A. J., Sutton, J., Harris, C. B., & Wilson, R. A. (2008). A conceptual and empirical framework for the social distribution of cognition: The case of memory. *Cognitive Systems Research*, 9, 33–51.
- Bengson, J., & Moffett, M. A. (Eds.). (2011). *Knowing how: Essays on knowledge, mind, and action*. New York: Oxford University Press.
- Bird, A. (2014). When is there a group that knows? Distributed cognition, scientific knowledge, and the social epistemic subject. In J. Lackey (Ed.), *Essays in collective epistemology* (pp. 42–63). Oxford: Oxford University Press.
- Blomberg, O. (2016). Common knowledge and reductionism about shared agency. *Australasian Journal of Philosophy*, 94(2), 315–326. <https://doi.org/10.1080/00048402.2015.1055581>.
- Bratman, M. E. (1992). Shared cooperative activity. *Philosophical Review*, 101, 327–341.
- Bratman, M. E. (1993). Shared intention. *Ethics*, 104, 97–113.
- Bratman, M. E. (2014). *Shared agency: A planning theory of acting together*. Oxford: Oxford University Press.
- Brownstein, M., & Michaelson, E. (2016). Doing without believing: Intellectualism, knowledge-how, and belief-attribution. *Synthese*, 193(9), 2815–2836. <https://doi.org/10.1007/s11229-015-0888-9>.
- Carter, J. A., & Pritchard, D. (2015). Knowledge-how and epistemic luck. *Noûs*, 49(3), 44–453. <https://doi.org/10.1111/nous.12054>.
- Cuijpers, L. S., Zaal, F. T. J. M., & de Poel, H. J. (2015). Rowing crew coordination dynamics at increasing stroke rates. *PLoS ONE*, 10(7), e0133527. <https://doi.org/10.1371/journal.pone.0133527>.
- de Brouwer, A. J., de Poel, H. J., & Hofmeister, M. J. (2013). Don't rock the boat: How antiphase crew coordination affects rowing. *PLoS ONE*, 8(1), e54996. <https://doi.org/10.1371/journal.pone.0054996>.
- de Poel, H. D., de Brouwer, A. J., & Cuijpers, L. S. (2016). Crew rowing: An archetype of interpersonal coordination dynamics. In P. Passos, K. Davids, & J. Y. Chow (Eds.), *Interpersonal coordination and performance in social systems* (pp. 150–153). London: Routledge.
- Feigean, M., R'Kiouak, M., Bootsma, R. J., & Bourbousson, J. (2017). Effects of intensive crew training on individual and collective characteristics of oar movement in rowing as a coxless pair. *Frontiers in Psychology*, 8, 1139. <https://doi.org/10.3389/fpsyg.2017.01139>.
- Fridland, E. (2013). Problems with intellectualism. *Philosophical Studies*, 165(3), 879–891.
- Fridland, E. (2014). They've lost control: Reflections on skill. *Synthese*, 191(12), 2729–2750.
- Fridland, E. (2015). Knowing how: Problems and considerations. *European Journal of Philosophy*, 23(3), 703–727. <https://doi.org/10.1111/ejop.12000>.
- Gilbert, M. (1987). Modelling collective belief. *Synthese*, 73, 185–204.
- Gilbert, M. (1989). *On social facts*. Princeton, NJ: Princeton University Press.
- Gilbert, M. (2000). *Sociality and responsibility: New essays in plural subject theory*. Lanham, MD: Rowman and Littlefield.
- Gilbert, M. (2005). Shared values, social unity, and liberty. *Public Affairs Quarterly*, 19(1), 25–49.
- Glick, E. (2015). Practical modes of presentation. *Noûs*, 49(3), 538–559.
- Goebel, W., & Palmer, C. (2009). Synchronization of timing and motion among performing musicians. *Music Perception*, 26(5), 427–438. <https://doi.org/10.1525/mp.2009.26.5.427>.
- Gold, N., & Sugden, R. (2007). Collective intentions and team agency. *Journal of Philosophy*, 104(3), 109–137.
- Harris, C. B., Barnier, A. J., Sutton, J., & Keil, P. G. (2014). Couples as socially distributed cognitive systems: Remembering in everyday social and material contexts. *Memory Studies*, 7(3), 285–287.
- Heyes, C. (2017). Apes submentalize. *Trends in Cognitive Sciences*, 21, 1–2.
- Knoblich, G., Butterfill, S., & Sebanz, N. (2011). Psychological research on joint action: Theory and data. In H. R. Brian (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 54, pp. 59–102). Amsterdam: Elsevier.
- Krupenye, C., Kano, F., Hirata, S., Call, J., & Tomasello, M. (2016). Great apes anticipate that other individuals will act according to false beliefs. *Science*, 354, 110–114.
- Lackey, J. (2014). Introduction. In J. Lackey (Ed.), *Essays in collective epistemology* (pp. 1–8). Oxford: Oxford University Press.

- Levinson, S. C. (1995). Interactional biases in human thinking. In E. N. Goody (Ed.), *Social intelligence and interaction* (pp. 221–260). Cambridge: Cambridge University Press.
- List, C., & Pettit, P. (2011). *Group agency: The possibility, design, and status of corporate agents*. New York: Oxford University Press.
- Mann, D. L., Spraford, W., & Abernethy, B. (2013). The head tracks and gaze predicts: How the world's best batters hit a ball. *PLoS ONE*, 8(3), e58289. <https://doi.org/10.1371/journal.pone.0058289>.
- Noë, A. (2005). Against intellectualism. *Analysis*, 65(288), 278–290.
- Palermos, S. O., & Tollefsen, D. P. (2018). Group know-how. In J. A. Carter, A. Clark, J. Kallestrup, S. O. Palermos, & D. Pritchard (Eds.), *Socially extended epistemology*. New York: Oxford University Press.
- Paternotte, C. (2013). The epistemic core of weak joint action. *Philosophical Psychology*, 28(1), 70–93. <https://doi.org/10.1080/09515089.2013.819614>.
- Pavese, C. (2015). Practical senses. *Philosopher's Imprint*, 15(29), 1–25.
- Pavese, C. (2016). “Skill in epistemology” parts I and II. *Philosophy Compass*, 11, 642–660.
- Pavese, C. (2017). Know how and gradability. *Philosophical Review*, 126(3), 345–383.
- Railton, P. (2009). Practical competence and fluent agency. In D. Sobel & S. Wall (Eds.), *Reasons for action* (pp. 81–115). Cambridge: Cambridge University Press.
- Santorio, P. (2016). Nonfactual know-how and the boundaries of semantics. *Philosophical Review*, 125(1), 35–82.
- Schweikard, D. P., & Schmid, H. B. (2013). Collective intentionality. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2013 edition), <http://plato.stanford.edu/archives/sum2013/entries/collective-intentionality/>. Accessed 5 Oct 2018.
- Sebanz, N., Bekkering, H., & Knoblich, G. (2006). Joint action: Bodies and minds moving together. *Trends in Cognitive Sciences*, 10(2), 70–76. <https://doi.org/10.1016/j.tics.2005.12.009>.
- Sebanz, N., & Knoblich, G. (2009). Prediction in joint action: What, when and where. *Topics in Cognitive Science*, 1(2), 353–367. <https://doi.org/10.1111/j.1756-8765.2009.01024.x>.
- Setiya, K. (2012). Knowing how. *Proceedings of the Aristotelian Society*, 62, 285–307.
- Skyrms, B. (2010). *Signals: Evolution, learning, and information*. Oxford: Oxford University Press.
- Stalnaker, R. (2012). Intellectualism and the objects of knowledge. *Philosophy and Phenomenological Research*, 85(3), 753–761.
- Stanley, J. (2011). *Know how*. Oxford: Oxford University Press.
- Stanley, J., & Williamson, T. (2001). Knowing how. *Journal of Philosophy*, 98, 411–444.
- Stanley, J., & Williamson, T. (2017). Skill. *Noûs*, 51(4), 713–726.
- Sterelny, K. (2012). *The evolved apprentice: How evolution made humans unique*. Cambridge, MA: MIT Press.
- Stix, G. (2014). The ‘It’ factor. *Scientific American*, 311, 72–79.
- Sutton, J., Harris, C. B., & Keil, P. G. (2010). The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive Sciences*, 9, 521–560.
- Tomasello, M. (2014). *A natural history of human thinking*. Cambridge, MA: Harvard University Press.
- Tomasello, M. (2016). *A natural history of human morality*. Cambridge, MA: Harvard University Press.
- Tuomela, R. (1995). *The importance of us—A study of basic social notions*. Stanford: Stanford University Press.
- Tuomela, R. (2007). *The philosophy of sociality: The shared point of view*. New York: Oxford University Press.
- Tuomela, R., & Miller, K. (1988). We-Intentions. *Philosophical Studies*, 53, 367–389.
- Vesper, C., Butterfill, S., Knoblich, G., & Sebanz, N. (2010). A minimal cognitive architecture for joint action. *Neural Networks*, 23, 998–1003.