

Knowledge Norms and Conversation

J. Adam Carter

Abstract: Might knowledge normatively govern *conversations* and not just their discrete constituent thoughts and (assertoric) actions? I answer yes, at least for a restricted class of conversations I call *aimed* conversations. On the view defended here, aimed conversations are governed by *participatory know-how* – viz., knowledge how to do what each interlocutor to the conversation shares a participatory intention to do by means of that conversation. In the specific case of conversations that are in the service of joint inquiry, the view defended is that interlocutors (A, B, ... n) must intentionally inquire together into whether p, by means of an aimed conversation X, only if A, B, ... n *know how* to use X to find out *together* whether p. The view is supported by considerations about instrumental rationality, shared intentionality, the epistemology of intentional action, as well as linguistic data.

§1.

A popular idea in epistemology holds that both thought and action are governed by epistemic norms, with *knowledge norms* being illustrative. Here are some popular contenders: knowledge norms on belief, assertion, and (non-assertoric) action.

Knowledge norm of belief: One must believe that p only if one knows that p.¹

Knowledge norm of assertion: One must assert that p only if one knows that p. (Williamson, 2000, 243)

Knowledge norm of action: One must treat the proposition that p as a reason for acting only if you know that p. (Hawthorne and Stanley, 2008, 578)

These (and other related) knowledge norms on thought and action are *distinct norms*, such that each governs a different type of thing, viz., a different thought or action type. For theoretical elegance, one might be inclined to accept all sorts of knowledge norms as part of a wider

¹ See Williamson (2000, 255–56). As Williamson puts it: “It is plausible (...) that occurrently believing p stands to asserting p as the inner stands to the outer. If so, the knowledge rule for assertion corresponds to the norm that one should believe p only if one knows p. Given that norm, it is not reasonable to believe p when one knows that one does not know p.”

‘knowledge-first’ approach²; even so, individual knowledge norms themselves (governing various types of thought and action) are typically argued for *independently*³, and what it is that each governs is taken to be independent from each other. If one believes that p, asserts that p, and then (non-assertorically) acts on the basis of p, in that order, then three tokenings are answerable to three separate knowledge norms, in that order.

My aim here is to consider how epistemic norms (with knowledge norms being the guiding illustrative focus), might, in some cases, apply to complex performances that involve dynamic combinations of thought and action.

I begin with a brief ‘proof of concept’ of this idea using practical reasoning as illustrative, and the main goal of this chapter will then be to then explore in some detail how *conversations*, or at least, a certain subset of conversations that are constitutively aimed, might themselves be knowledge governed. On the proposal I will develop, knowledge norms are applicable to aimed conversations that involve thinking *to a purpose*, practical or theoretical; aimed conversations are governed by *participatory know-how* – viz., know how to do what each interlocutor to the conversation shares a participatory intention to do. When individuals intentionally inquire together into whether p, by means of an aimed conversation X, they ought to know how to use X to find out together whether p. While participatory knowledge-how is in the above respects the norm of (aimed) conversation, satisfying it will in practice generally require interlocutors to the conversation thinking and acting in accordance with more specific knowledge norms (e.g., of assertion and inquiry). But aimed conversations are governed by know-how in a way that is not *reducible* to any concatenation of satisfying individual norms on thought and action.

§2.

A key idea in Williamson’s (2017) picture of the relationship between epistemology and action theory is that there are important structural connections (and associated analogies) between thought and action, and between the norms that govern them.

My aim here isn’t to evaluate these structural analogies.⁴ But rather to take Williamson’s idea of the *operation of a cognitive-practical system* as illustrative of the kind of thing that can (*apart* from, and in addition to, the discrete elements of its operation) be norm governed.

² See Williamson (2013) for an overview of this approach. See also Ichikawa and Jenkins (2017) for some helpful discussion of how ‘knowledge-first’ has been associated with a cluster of claims rather than with any single claim.

³ This is the case even when these arguments pursue similar strategy types, e.g., retraction data, lottery cases, Moore-paradoxical assertions, etc.

⁴ Though see Miracchi and Carter (2022).

A cognitive-practical system, such as human thinkers who manifest intelligence paradigmatically by reasoning and then behaving on the basis of that reasoning (*viz.*, by representing the world, then purposefully changing it), exhibit two dimensions of intelligence with opposite directions of fit, *cognitive* and *practical*, where “cognitive concerns those aspects of intelligent life which concern fitting mind to world, ‘practical’ those which concern fitting world to mind”. (Williamson, 2017, 164)

On this picture, *practical reasoning* is, as Williamson puts it, the ‘nexus’ between the cognitive and the practical (between the mind-to-world and world-to-mind dimensions of intelligent life). On the cognitive side, we have beliefs, and, if all goes well, knowledge (the ‘gold standard’ in fitting mind to world). Representational attitudes such as belief and knowledge serve as *inputs* to practical reasoning, reasoning about what to do. The output of practical reasoning (which will have a world-to-mind direction of fit) isn’t just another belief, but an *intention*⁵ to do something, and all goes well here only if one *acts* (and not merely *intends* but fails to act).

Importantly though, a *cognitive-practical* system like us isn’t, in practical reasoning, *merely* doing something that involves a mind-to-world direction of fit, or merely doing something that involves a world-to-mind direction of fit. Rather, it’s – in *practical* reasoning – doing *both* kinds of things, and not in isolation from each other, but in such a way that world-to-mind intelligence is supposed to *inform* mind-to-world intelligence. As such, the operation of a practical reasoning system can fall short on *either* side, e.g., if one has an input short of knowledge or an output short of action. Here’s Williamson (2017, 174):

When all goes well in one’s whole cognitive-practical system (not just one’s brain), one acts on what one knows. But things often go badly. One believes that P when not P, or (perhaps as a consequence) intends to ϕ but never ϕ s. In practical reasoning, the false belief plays the same local role as knowledge (it makes a premise), and the failed intention plays the same local role as action (it makes the conclusion), but they do not play the same global roles. Such malfunctions have to be understood in relation to what happens when the whole system is functioning well. (...) When the cognitive-practical system functions well, one acts on what one know.

While there are both belief-like and actional components of the operation of a cognitive-practical system, the *operation of the system*, *viz.*, practical reasoning, itself involves a complex consisting of the beliefs and the (belief-informed) actional output. And this complex operation is itself norm-governed on Williamson’s approach by a *complex* knowledge-action norm, which governs the operation of a cognitive-practical system:

⁵ This is the position of the updated (Williamson, 2017) analogy. See Williamson (2000, Ch. 1) for a previous statement of the analogy which holds that belief is to knowledge as desire is to action.

Knowledge-action norm on practical reasoning (KAN): One must: reason practically only if one acts on what one knows.

A quick point of clarification. There is a crucial difference between the above ‘knowledge-action’ norm on practical reasoning, and what is sometimes called the ‘knowledge norm of action’ or alternatively the ‘knowledge norm of practical reasoning’ as defended by Hawthorne and Stanley (2008):

Knowledge norm of action: One must: treat the proposition that p as a reason for acting only if one knows that p.⁶

These norms govern different types of things. Hawthorne and Stanley’s knowledge norm governs *reasons treating* – viz., it governs *treating propositions as a reason to act, viz.,* – our appropriation of that *for which* we act. And possessing knowledge is the standard for doing *that*.

By contrast, KAN says knowledge *and* action, and the latter being done on the former, constitute the satisfaction of the (complex) norm for doing something else. The something else is the activity of *practical reasoning* itself.

The point of this section has been to clarify the very idea of what it would be for a knowledge-involving norm (with KAN being our example) to govern *complexes of thought and action*, even when we think that the constituent thought and action tokens that feature in the complex might be individually governed themselves by knowledge norms on thought and action types as such. Williamson’s view of the operation of a cognitive practical system – practical reasoning performances – offers a precedent here.

With this precedent in view, I now want explore in what follows the idea that a *conversation* can be modelled in a way (broadly analogously to the way Williamson has modelled the operation of a practical reasoning system) such that it is plausibly constrained by a complex knowledge-implicating norm, over and above whatever knowledge norms might govern the constituent thoughts and actions that feature in a given conversation. §3 will take this up, beginning with several up-front clarifications and comments on scope.

§3.

What kind of thing is a conversation? Our colloquial usage is liberal here as to what would qualify. Likewise, philosophical delineations of conversations err on the side of inclusiveness. Grice’s (1975) theory of conversational implicature, for example, takes a ‘conversation’ to be

⁶ See Hawthorne and Stanley (2008, 578).

minimally “a communicative exchange”, which will (typically) at least involve two interlocutors, and where at least some information is communicated through speech acts by those interlocutors.

A simple example of a conversation between interlocutors A and B might feature just two speech acts: a question (A) and an answer (B). Without knowing what that question was, or what the answer was, there is already a *trivial* sense in which a simple conversation with that {Q +A} structure might be governed by knowledge norms: one per speech act.

B’s answer (whatever it was), communicated through assertion, would (on a knowledge norm of assertion) be proper only if B *knows* the answer to A’s question, whatever it was. What about A’s question? A recent idea, defended by Christopher Willard-Kyle (forthcoming) is that inquiring via interrogative attitudes also requires (a specific kind of) knowledge in order to be proper as an instance of inquiry, and this is the case even if we grant that we would typically have a reason to ask a question only if we were ignorant of what the answer was. On Willard-Kyle’s view, there nonetheless is something defective about inquiring into something you don’t *know* has a true (and direct) answer. Consider the following:

- 1) Why does the sun set in the east?
- 2) Was it Abraham Lincoln or Franklin Delano Roosevelt who was the first U.S. president?⁷
- 3) Is a person with 32 hairs bald?

Each of 1-3 is a defective *as* a question. (1) and (2) don’t have true direct answers. And (3) is a defective question even if it did have a true direct answer so long as the individual asking this question doesn’t know, when asking it, *that* it has a true direct answer. The formulation of the knowledge norm for inquiry Willard-Kyle defends, and which gets these results, holds that:

Knowledge Norm for Inquiry: One ought to: inquire into (an unconditional question) Q at t only if one knows at t that Q has a true (complete, and direct) answer.

We can set aside for now what counts as a complete, true and direct answer (there is some debate here⁸), and we can also set aside whether Willard-Kyle’s preferred norm is correct. While I take it to be plausible, I won’t try to defend it here.

⁷ These examples in (1) and (2) are due to Willard-Kyle (forthcoming).

⁸ For some discussion here, see Willard-Kyle (2023).

Rather, the point is just that, on a toy conversation that involves a single {Q+A} pair, we can see how that conversation would (in a trivial way) be governed by knowledge norms *in so far as* the conversation has two components (a question and an answer), where (i) the question beginning the conversation is governed by the knowledge norm for inquiry, and (ii) the answer ending the conversation is governed by the knowledge norm of assertion.

But notice that *this* sense in which a conversation might be governed by knowledge-implicating norms is not one such that the *complex* communicative exchange, qua the kind of thing *it* is (*viz.*, a *conversation*), is epistemically norm governed over and above the simple *aggregation* of the epistemic norms that govern the discrete speech acts that make the conversation up.

Even if we suppose the individual speech act components of a conversation (in the {Q+A} example above) are governed, respectively, by the knowledge norm of inquiry and the knowledge norm of assertion, is there any reason to think that conversations, *as such*, might answer to any kind of more complex, knowledge implicating norm?

Remember, conversations (taking the minimal Gricean idea as illustrative) are in essence just “communicative exchanges”. The below are all conversations:

Conversation 1 A: “France is a country.” B: “I am hungry” A: “Spain is a country.”

Conversation 2 A: “Huh.” B: “Eh?”

Conversation 3 A: “How’s it going?” B: “How’s it going?”

These are all conversations in the sense that they are communicative exchanges. None is epistemically impressive. Is any defective *qua* conversation in a way such that, e.g., any kind of candidate knowledge norm governing conversations as such would capture this intuitive defectiveness *independent* of the aggregate sense in which the constituent speech acts might be separately knowledge governed?

We might be tempted to think so. Maybe we think that, in Conversation 1, B, simply in virtue of participating in this conversation, should have known to talk about France (or at least in some way acknowledge A’s claim about it); or, perhaps A should have known to address B’s hunger rather than talking about Spain.

But we actually have too little information to make this kind of assessment. For all we know, the first conversation might have occurred against a background where there was an agreement that A would continue telling B what the countries are up until the point where B mentioned their hunger twice. Similarly, for all we know, in Conversation 2, A was expressing an attitude, and B was querying the context. Without additional information about there being any *further aim* to this communicative exchange over and above the individual aims constituent in each speech acts

(an attitude expression, and a question), it's not clear how there would be any overarching knowledge-implicating norm of conversation violated here. The same goes for Conversation 3. Colloquially, "How's it going" is the use of a sentence in the form of a question in order to perform the act of *greeting* someone. Conversation 3 is plausibly just a matter of two individual acts of greeting, the second in response to the first – which very plausibly might have been the (exhaustive) purpose and extent of A and B communicating at all here.

What the brief look at minimal and *prima facie* epistemically unimpressive conversations 1-3 suggests is this: without clear information about what a conversation's *purpose* is, (regardless of how such a purpose would be established – this is some-thing we'll take up in the next section), we aren't in a position to work out whether any kind of knowledge-implicating norm would potentially govern particular communicative exchanges *as a conversation*. The most we can do (absent such information about conversational purpose) is normatively assess the individual speech acts featuring in the conversation relative to norms that might govern those specific acts (e.g., the knowledge norm of assertion, the knowledge norm of inquiry, etc.). And this is so even though these conversations seem to be unimpressive epistemically.

§4.

Call a conversation an *aimed conversation* where the participants in the conversation have a goal that they intentionally pursue, together, by means of speech acts. Paradigmatic aimed conversations are those where interlocutors are *thinking to a purpose*, such as the purpose of, together, *figuring something out*. A conversation with this aim would involve thinking to a purpose insofar as it involves thinking to a *theoretical* purpose: a joint *inquiry*. A toy example of an aimed conversation which involves thinking to a theoretical purpose is a jury deciding together one question: whether p.

A clarification: not all joint inquiries are pursued by means of conversation, at least not in any colloquial sense. Individuals might inquire together, e.g., in cases of *distributed cognition*, where individual tasks are isolated. (Consider here Edwin Hutchins' (1995) example of a ship crew navigating to port, figuring out where to steer the ship, each by doing separate things). So joint inquiry is a wider category than aimed conversation that involves thinking to a (theoretical) purpose.

Another kind of *aimed conversation* involves thinking to a purpose that is *practical*. Here participants in the conversation have a practical goal that they *indirectly* pursue, intentionally, by means of speech acts. Such conversations function as (at least part of) the 'coordination' of coordinated action. A toy example of an aimed conversation that involves thinking to a practical purpose is a conversation in which interlocutors are communicating in order to get medical help. The goal of the conversation might be that of *getting help*.

It is *aimed conversations* that I'll defend as plausibly norm governed. Before looking squarely at what such norms would be, let's pause to clarify in some more detail what would plausibly distinguish aimed conversations (in the sense I'll be interested in, and which will matter for their being norm-governed) from conversations like 1-3 above.

There are two dimensions of such aimed conversations I want to highlight. One concerns the *intentional* character of the conversation, toward its aim, and the other concerns the *shared* dimension by which the aim is pursued by interlocutors through speech acts by which they are thinking, in aimed conversations, *together* to a purpose.

First, regarding the *intentionality* of aimed conversations. Compare the following –first, an unaimed conversation, consisting in a communicative exchange between two individuals, A and B, while waiting in a queue. They might discuss a range of topics. The speech acts through which they think together about these topics might themselves be intentional speech acts. *But*, suppose A is asked afterwards to describe the conversation had with B: A might do this by recollecting how the conversation went in much the same way as a third-party observer might do this. Relatedly: suppose prior to joining the queue and seeing B standing there, A recalls how conversations often go with B and predicts that the conversation will likely be about the weekend's gossip (which it is).

One of the epistemic hallmarks of intentional actions is that the knowledge we have of them in action is, in some important respect, *non-observational*. A famous example of Anscombe's here involves intentionally opening a window. If I intentionally open a window and in doing so happen to cast a spot of light on the wall, I know about the window-opening in a way that is not merely observational, which is how I know about the casting of light on the wall.⁹ The kind of aimed conversations I am interested in are ones we'd know about, *as participants thinking together*, in a way that is like how we'd know we're opening a window when we are and not how like we know it casts a spot on the wall when it does. Secondly, on the non (merely) predictive character of intentional actions. In intentionally ϕ ing by means ψ (say, hailing a cab by raising a hand) our knowledge of the relationship between the means and end is not merely predictive, as it might be if you observed someone raising a hand followed by a taxi approaching. In intentionally hailing a cab, you have a kind of knowledge with *means-end structure*; you know that you are ψ -ing *in order to* ϕ .¹⁰

Contrast now a simple jury deliberation about whether p, with our water-cooler conversation described above. The latter, we've seen, is such that our knowledge of how the conversation went would be largely observational, and if it ended up going a particular way (reaching some

⁹ For discussion of this example in relation to non-observational knowledge of intentional action, see Schwenkler (2012).

¹⁰ For discussion of how acting intentionally characteristically involves knowledge with means-end structure, see, e.g., Anscombe (1957, §§28-9); Paul (2009, 5); Mele and Moser (1994, 63) and Pavese (2022, 375).

particular conversational conclusion), we might have predicted that, but we did not at the time know we were making conversational moves *in order to reach that point* whatsoever.

To this point I've suggested *aimed conversations* are intentional in a way that lines up with epistemic features of intentional actions. But even if the individuals each perform speech acts *intentionally* in an aimed conversation, there is also a sense in which aimed conversations (e.g., between two interlocutors) involve *shared intention* – and not merely shared in the sense that the content of each of the individual intentions of each interlocutor is the same.

The idea of a shared intention is central to most conceptions of *coordinated* activity, though its nature is a point of some theoretical debate.¹¹ For my purposes here, I'll at least note a few of the well-known characteristic *functional* features of shared intentions (setting aside the comparatively thornier matter of how exactly to sub-stantively characterise the nature of such intentions), as developed influentially by Bratman (2013), and register what kind of shape these functional features take in aimed conversations specifically.

On Bratman's view, when two people (A and B) do something ϕ together, with a *shared intention*, each individual has a *participatory intention* that takes the form (on the part of A), that *A and B ϕ* . So if we take a walk together, we each intend that *we* take a walk. And our participatory intentions then structure the way we coordinate activities, plan, and bargain. To give a simple example, if A and B have a shared intention that A and B walk together, and B needs to tie B's shoe, A will stop too and wait. And A's stopping is in order that A and B walk together, etc. Likewise, if A has a strong preference which way to go on the walk, B will (*ceteris paribus*) agree to go that way (even if doing so is not B's individual preference apart from the participatory intention) in order for both A and B to walk together, perhaps while also expecting that A would do the same for B, for the same reason. In this way, A's and B's shared intentions structure what A and B do (in terms of executing subplans related to the shared intention), how they bargain with each other, when they take a walk *together*, with the shared intention that they do this, rather than just walk individually near each other.

What goes for having a shared intention to do something, ϕ , together, goes like-wise for a shared intention to *have a conversation together*. Here though a scope clarification is important. Consider two cases in which A and B share an intention to have a conversation together, and that this intention is a participatory intention in Bratman's sense. It might be that the *content* of the shared intention here is just that *we have a conversation* – *viz.*, that, for some conversation, X, we have X. This needn't involve any kind of thinking to any particular purpose, theoretical or practical. Compare this again with our toy example of an aimed conversation, where we (together) are trying to *figure something out*, and we use a conversation as a means towards this end. Our participatory intention here will not be that we have *any*

¹¹ See, along with Bratman (2013) also, e.g., Gilbert (2009; 2008) and Velleman (1997).

old conversation; it is that we *figure out whether p, by means of conversation* – viz., that we have a conversation that achieves that result.

Putting the ideas in this section together, we have now got a grip on some of the key features of *aimed conversations* of the sort that will be of special interest in so far as we can construe such conversations *themselves* (and not just their individual speech act components) as governed by knowledge-implicating norms. In summary: *aimed conversations* are conversations (minimally, communicative exchanges) where participants in the conversation have a goal that they intentionally pursue, together, by means of speech acts. The goal might be theoretical or practical. Participants *intentionally* pursue the goal together in so far as they have participatory intentions that structure their behaviour (e.g., subplans, bargaining) towards that goal. And the *actions* each takes in using speech acts toward bringing about what they share a participatory intention to bring about through conversation (e.g., when thinking to a theoretical purpose, to figure out together whether p), are actions each stands to *epistemically* in a way that lines up with our epistemic position towards our intentional actions generally, in so far as this knowledge is non-observational and known by participants to have a *means-end* structure.

When conversations have the above features, and in particular – when participants have a shared intention to think together (by means of a conversation) to a particular purpose (e.g., settling a whether p question), the purpose toward which they think together intentionally is partly *constitutive* of the kind of thing they are doing by means of speech acts; their conversation is *constitutively* aimed. That the participants are using speech acts to think together toward *that* purpose distinguishes what it is they are doing together from other things they might do together with speech acts, such as have a conversation *for conversation's sake*, or, e.g., for the sake of convincing passers-by they are on friendly terms.

When we are engaged in constitutively aimed activities (chess, baseball, etc.) and someone performs in a way that doesn't accord with the rules of the game, it is appropriate to *point this out*. For instance, if we are playing baseball, and a teammate lies down in the outfield to count the grass, we might say "What are you *doing*?" If they reply and tell you what they're doing (counting grass), you would appropriately re-join with, "But that's now what *we're doing* – *we're* playing baseball!" The appropriateness of this kind of rebuke is explained by the fact that the shared intention of *playing baseball* together implicates the kinds of subplans, bargaining, etc., that are bound up with any genuine kind of participatory intentions, where counting grass doesn't feature in any plausible associated subplans. And what goes for constitutively aimed activities like baseball goes for constitutively (theoretically) aimed conversations, where interlocutors' participatory intentions are (by means of speech acts) to *figure out whether p*.

Compare here again with the baseball example: if you and I are trying to figure out whether either of us has ever travelled to a country with a non-extradition agreement with the United States, there are several things we need to figure out together in our communicative exchange. We need to figure out

which countries actually *lack* extradition agreements with the U.S. We begin alphabetically, and work out that such countries at least include Afghanistan, Algeria, Andorra, Angola, Armenia, Azerbaijan, and Bahrain. We've not been to these, suppose, so we keep going, still unsure if either of us has been to such a country. Suppose that during this process (as we work down the alphabet), you begin telling me facts about Bahrain, given that we have after all just brought it up. I listen politely for a bit, and I'm learning a lot, e.g., that until the 3rd century BC, it was part of the Achaemenid Empire. Although you're telling me things I'm interested in, and which you know, and I'm enjoying this, and you're enjoying this, your speech acts are still criticisable in the same way as, when playing baseball, we might ask the person lying on the ground "what are you *doing*?" Here, I might say also say, "what are you *doing*? *We're* trying to figure out whether we've been to a non-extradition country with the U.S."

Notice that the above rebuke is permissible even when your speech acts, and my speech acts, satisfy knowledge norms individually. That is, even when you know what you say about Bahrain (satisfying the knowledge norm of assertion), and even if I join you in straying from what we're doing (inquiring about U.S. non-extradition countries) by asking follow-ups about Bahrain, while knowing there is a true, complete answer to my question, etc.

With a view now as to how (constitutively) aimed conversations¹² can be norm governed *apart* from the individual speech acts that feature in them, let's turn now to a consideration of how such norms governing these conversations might be, specifically, *knowledge* norms.

§5.

When two people are engaged in a (constitutively) aimed conversation to figure out whether p, they are *inquiring* (into a question) by means of conversation. It is worth registering a point of connection between norms on *inquiry* generally, and the kind norm we'd be in the market to capture on an aimed conversation.

As Jane Friedman (2020) points out, once it's settled that what one is doing is inquiring into whether p, viz., that it is one's aim to figure out whether p, one is thereby committed, by reference to a norm of *instrumental rationality*, to take the necessary means to figuring out what one is inquiring into. Friedman calls this a *zetetic instrumental principle (ZIP)*:

Zetetic Instrumental Principle (ZIP): If one wants to figure out Q, then one ought to take the necessary means to figuring out Q.

¹² That is, of the sort outlined in this section, viz., as when interlocutors think together to a (theoretical) purpose.

The idea developed in the previous section bears some similarities to ZIP. Consider the following principle, which codifies the idea captured by the baseball/aimed conversation analogy:

Constitutive aim/Instrumental Rationality (CAIR): When some activity one engages in intentionally has a constitutive aim, then one ought to take the necessary means to bringing about that aim.

When one intentionally goes bowling, one ought to throw at the pin.¹³ The ‘ought’ here that applies to what one ought to do when engaged in an activity that is constitutively aim governed is an *aim relative* ought in the sense that Friedman’s ZIP norm is also an instance of an aim-relative ought. Friedman’s aim-relative norm on inquiry cashes out the ‘aim’ in terms of what one *wants* (i.e., where the aim is specifically desire-relative). CAIR is a bit different here; it gives a conditional where what suffices to make it such that one ought to take the necessary means to bringing about a given aim *is that one is intentionally engaging in an activity constitutively governed by that aim*. That, on CAIR, is taken to be sufficient for being such that one ought to take the necessary means to bring about that aim. And to reiterate the ‘ought’ (as with ZIP) is with CAIR not an all-things considered ought, but an aim-relative ought.

CAIR is not about conversations *per se*. The special case of CAIR that will be relevant for generating knowledge norms on (constitutively) aimed *conversations* (with theoretical aims) will be more specific. A first general statement is:

Constitutive aim/Instrumental Rationality – Aimed Conversations (CAIR-AC): When individuals (A, B, ... n) intentionally inquire together into whether p, by means of an aimed conversation, then (A, B, ... n) ought to take the necessary means to figuring out whether p.

Since aimed conversations, in the sense we’ve characterized already, involve *shared intentions*, we can unpack CAIR-AC more fully, making explicit the connection between shared and participatory intentions:

Constitutive aim/Instrumental Rationality – Aimed Conversations (CAIR-AC*): When individuals (A, B, ... n) intentionally inquire together into whether p, by means of an aimed conversation, then (A,

¹³ Note that there are stronger versions of this norm that don’t require the ‘intentional’ caveat. For example, Michael Lynch (2009) thinks that simply by virtue of participating in any activity that is constitutively norms governed, one is thereby committed to pursuing the value of the relevant activity’s aim. This, he takes it, explains why simply being involved in the activity of inquiry commits one to the value of truth and to pursuing it.

B, ... n) ought to take the necessary conversational means to figuring out whether p, and these means will include those structured by the *participatory intentions* (A, B, ... n) share.

Recall that when we have *participatory intentions* (e.g., I intend that *we* find out whether p, and you intend that *we* find out whether p), these structure how we negotiate, form subplans, bargain, etc., accordingly.

It is not hard to see how the requirement that participants intentionally inquiring together into whether p (by means of an aimed conversation) satisfy *individual* knowledge norms governing speech acts used (as a means to figuring out whether p) will fall out of CAIR-AC*. Consider the following further twist on our previous example case, where we are having a conversation to figure out who has been to a country without an extradition agreement with the US. Now, suppose the conversation includes this exchange:

A: "Tajikistan is a U.S. non-extradition country, and I've not been there."

B: "Neither have I. OK, so we still don't know if either of us has been to one. Let's think further. Have you been to Canada?"

A: "Yes."

Note, for one thing, that interlocutors who fail a *knowledge norm of assertion* will inevitably also fall short by the lights of CAIR-AC*. Tajikistan is a U.S. non- extradition country. Suppose A is confused and *has* in fact been to Tajikistan.

On this assumption, A asserts what A doesn't know because it's false. Making such an assertion is "not" taking necessary conversational means to figuring out together whether p; on the contrary, it would be contravening the goal of figuring out together whether p (whether either has been to such a country). On the contrary, if A asserts "Tajikistan is a U.S. non-extradition country, and I've not been there" knowledgeably, that constitutes taking means necessary (working through these countries together until they figure out whether either has been to one) to figuring out whether p. So A's first assertion "Tajikistan is a U.S. non-extradition country, and I've not been there" looks like one that would constitute a speech act in alignment with CAIR-AC* only if the act itself meets the knowledge norm of assertion.

The analogous point holds with respect to a Willard-Kyle-style *knowledge norm of questioning* – though this point requires some delicacy. When B asks "Have you been to Canada?", B will of course know that this question (framed as such) has a true, direct answer. *If Canada* was not even a country, B would, through asking this question, fail the knowledge norm of questioning and by extension also (by asking a question without a true,

direct answer, as would be the case if Canada wasn't a country) fail to be contributing through this speech act to the shared goal of figuring out p, as per (CAIR-AC*).

Importantly, though, in the above conversation, we know that B is (by asking *that* question) clearly *failing* CAIR-AC*, *even though* Canada is a country. So why is this? And does this mean that for one's speech acts to satisfy (CAIR-AC*) it's not enough that they satisfy knowledge norms governing individual speech acts?

The reason that B's question "Have you been to Canada" fails (CAIR-AC*) is that, since Canada *does* have an extradition treaty with the U.S., getting an answer to this question is *useless* towards the aim of figuring out the target question that A and B are pursuing. One might think though that the question asked by use of the sentence "Have you been to Canada" is actually, given the conversational context in which the sentence is used, a question expressing a false presupposition (and thus one that lacks a direct, true, answer) – namely, the question *Have you been to Canada, which is a U.S. non-extradition country?* That qualified question, if asked, actually fails the knowledge norm of inquiry given that that question lacks a true, direct answer, on account of having a false (semantic) presupposition. If the latter is the right way to think about the question asked by B about Canada, then it might seem that what goes wrong when one's speech act fails (CAIR-AC*) is going to be determined entirely by whether one satisfies *individual* knowledge norms on one's speech acts – e.g., on questions and answers.

This, however, is mistaken. To see why, we can easily imagine conversations with questions and assertions that are impeccable by the lights of knowledge norms on those particular speech act types, but which fail CAIR-AC* straightforwardly. In fact, we need just run a twist on our previous sample conversation (supposing A knows what A asserts), and change B's final question to: "Have you been to Canada, *which I know is not a US non-extradition country?*" This question is one that B *knows* admits of a true, direct answer (it lacks any false presupposition – that would be cancelled here by the added qualification) – and yet this speech act is *not* (by CAIR-AC*) actually an instance of taking necessary conversational means to figuring out whether p. It's at best chasing a red herring or stalling, at *worst* (assuming A's been to Canada) contributing to getting the wrong answer to the inquiry they're pursuing together.¹⁴

§6.

It might seem that the sense in which satisfying CAIR-AC* requires satisfying knowledge norms could be specified as a conjunction of satisfying (i) knowledge norms on particular speech acts used (e.g., assertion, questioning) in the aimed conversation; and (ii) Gricean

¹⁴ At least, it would be the wrong result if neither has in fact been to a U.S. non-extradition country. Then by thinking they had in virtue of one of them traveling to Canada, they'd conclude their inquiry incorrectly.

cooperative norms of quantity, quality, relation, and manner. B's question "Have you been to Canada, *which I know is not a US non-extradition country?*", in the context of the conversation discussed above, satisfies the knowledge norm of inquiry while failing to be cooperative via the cooperative norm on *relation*. It is irrelevant to the conversation. So: will individuals intentionally inquiring together into whether p, by means of an aimed conversation, CAIR-AC*, take the necessary conversational means to figuring out whether p *so long as they satisfy (i) and (ii)*, viz., knowledge norms on speech acts and Gricean norms pertinent to the conversation?

The answer is – for several reasons – ‘no’, and seeing why will be instructive for getting a clearer picture of what kind of knowledge norm is best understood as governing such aimed conversations by means of which individuals inquire together.

Firstly, notice that being cooperative (as per Grice) *conditioned upon* not knowing what is in the common ground is a recipe for failing to take the necessary conversational means to figuring out whether p. Secondly, one could be well intended to be cooperative but *lack the knowledge* for how to do so, in virtue of failing to know *how* the use of speech acts such as assertions or questions could most expediently bring about the aim of figuring out *together* (with one's interlocutor) whether p.

Consider here, by way of comparison, Williamson's view of what's required for a cognitive-practical system to operate well; the gold standard here is *acting on knowledge*; mere lucky true beliefs can sometimes play the same local role in a cognitive practical system as knowledge, and yet the system as a whole still isn't operating as it should. One is acting on less than knowledge.

Likewise, consider that *even if one didn't know how to use* speech acts to bring about their shared aim, one might luckily, along with one's interlocutor, stumble upon the aim of the joint inquiry. Consider the following example inquiry (suppose A and B are pursuing this together):

Married-Unmarried Riddle: *Alice, Bob and Charles are standing in a line. Alice is looking at Bob. Bob is looking at Charles. We know that Alice is married, while Charles is not. Is there a married person among the three who is looking at an unmarried person?*¹⁵

A and B are initially flummoxed. They explore (through questions and assertions) whether there is even enough information here to figure out the answer. Then A asks who Charles is looking at. B asks whether Bob is married to either Alice or Charles. These questions are all false leads, in that their answers would be use- less to finding out the answer. Suppose a third interlocutor, "C", supplies A and B with a hint: "It doesn't matter whether

¹⁵ See <https://codedrills.io/contests/icpc-amritapuri-2020-preliminary-round/problems/the-married-unmarried-riddle#>

Bob is married or unmarried.” A and B don’t know whether to trust C here, but they try this hint out and through a conversational back and forth stumble upon the answer – if Bob is married, he’s looking at an unmarried person, Charles. If he’s unmarried, a married person (Alice) is looking at him.

A and B didn’t appreciate why this hint was helpful, but just blindly followed the lead. In this case, there is a sense in which even when (following the prompt) A and B begin asking questions and making assertions that not only themselves follow knowledge norms, *and* are following Gricean maxims (their answers are all cooperative/useful) they *together* are still not functioning as they should, in a way that is analogous to how a cognitive-practical system is not functioning as it should when acting on merely lucky true beliefs. Similarly, A and B didn’t suffice to (as per CAIR-AC*) take the necessary conversational means to figuring out whether p: rather, it was A, B, *and* C – A and B’s conversational means employed were, absent C’s contribution, insufficient, despite their participatory intentions, to figure this out *together*.

Moreover, notice that in the above example, A and B fall short of satisfying CAIR-AC* for an important and easily overlooked reason. They are *intentionally* inquiring together into whether p.

§7.

A hallmark of action theory holds that even when one intends to ϕ , and then ϕ s, it’s still an open question whether their ϕ ing itself qualifies as intentional; in order to do that, viz., in order to *intentionally* ϕ , one’s success in ϕ ing must be controlled in such a way as to be suitably non-accidental. For example, I can form an intention to win an art contest, paint what is an average painting, get lucky and in fact win that art contest. But I didn’t *intentionally win the art contest*. At most I intentionally painted, and submitted the art.

One idea of Anscombe’s, originally dismissed due to counterexamples by Davidson (1970)¹⁶, but which is regaining popularity, is that doing something intentionally requires knowing *that* you are doing it. This would explain why I didn’t intentionally win the art contest; I knew that I was entering, but nothing I did was such that I knew I was winning when I did it. On this Anscombe-inspired way of thinking, acting intentionally is luck or accident precluding on account of the fact that acting intentionally requires acting with *propositional knowledge* and not merely with lucky belief.¹⁷

There’s considerable debate in recent action theory about whether Anscombe’s propositional knowledge condition on intentional action (or some version of it) is plausible, or whether this condition is perhaps too strong.¹⁸ Even so, most all sides of these debates are sympathetic to a

¹⁶ Though for a recent causalist approach to (luck precluding) intentional action, see Kelley (2022).

¹⁷ For discussion, see Kearl (2023).

¹⁸ For some recent discussion here, see Piñeros-Glasscock (2020); Beddor and Pavese

different kind of knowledge condition on intentional action, one that requires at least *know-how*:

Knowledge-How/Intentionality (KHI): If *S* intentionally ϕ s, *S* knows how to ϕ .¹⁹

You don't intentionally pick a lock unless you know *how* to pick the lock. You don't intentionally win the lottery *because* you don't know how to actually do it, you just know how to take a means by which you might then get very lucky and win.

Of course, there is dispute over how to characterize the nature of knowledge how, whether it is propositional or dispositional – at least, that is the way that the debate is often framed. However, for the purpose I'll now explore, this contrast is (as intellectualists about know-how have pointed out) crucially misleading for the following reason: *even if* knowing how to do something is reducible to *knowing a proposition* (e.g., on the Stanley and Williamson (2001) line – this will be knowing, of some way, *w*, for you to ϕ , that *w* is a way for you to ϕ), the way one has to know this proposition in order to count as knowing *how* to ϕ is special: one must know this proposition under what's called a *practical mode of presentation*.²⁰ To use a simple example here, consider one's knowledge how to use a pair of pliers. In order for one to know *how* to use the pair of pliers (as opposed to merely knowing how *one* uses pliers), one must be able to *think of using* the pliers as a way for oneself to use them – and in thinking this way you will be essentially disposed to, e.g., reach for them in a certain way, manipulate the handle in a certain way, to a certain purpose, etc.²¹

The matter of how to characterize practical modes of presentation within an intellectualist theory of know-how is itself a matter of dispute, though we needn't delve into this here. What matters for the present purposes is simply this: (i) KHI is almost universally held, and (ii) however one characterizes the nature of know-how in KHI (i.e., regardless of whether propositional attitudes or success-oriented dispositions are taken to be comparatively more *fundamental* in an account of knowing how to do something), we should expect that one will countenance (in one way or another) the importance of success-oriented dispositions to know-how possession.

These points give us a useful perspective to think about how *aimed conversations* themselves (regardless of how complex these are), and not just their discrete speech act components, might plausibly be governed by knowledge norms, or at least, a certain kind of knowledge norm.

(2021); Paul (2009); Carter and Shepherd (2023); Shepherd and Carter (Forthcoming).

¹⁹ See Pavese (2021, §5).

²⁰ Or, as Pavese (2015) characterizes it, a 'practical sense'.

²¹ See also Stanley (2011) for an expansion of this idea.

The idea, to a first approximation, is that joint inquiry, by means of aimed conversation, is going as it should when and only when interlocutors *knowhow* to do what each has a participatory intention to do. In slogan form: *participatory know-how is the norm of aimed conversation*; and in the special case of joint inquiry, interlocutors must inquire together through aimed conversation (to find out whether p) only if they actually know how to use speech acts effectively in the service of togetherbringing about the shared theoretical goal. That is:

Aimed conversations/know-how (ACK-1): Aimed conversations are governed by participatory *know-how*.

Aimed conversations/know-how (ACK-2): Individuals (A, B, ... n) must: intentionally inquire together into whether p, by means of an aimed conversation X, only if A,B,...n know how to use X to find out together whether p.

ACK-1, as a general thesis, is motivated by putting KHI together with a feature of aimed conversations, which is that their (constitutive) aim (e.g., to figure out whether something is so) is grounded in interlocutors' shared participatory intentions. ACK-2 spells out the content of ACK-1 by saying *how* aimed conversations are governed by participatory know-how in the special case of interested to us – joint inquiry by means of conversation – where the aimed conversation has the theoretical aim of (interlocutors) figuring out together whether p.

Along with the above principled motivation for ACK-1 and ACK-2, there are also familiar linguistic data points that speak in favour of participatory know-how as governing aimed conversation. I will mention a few and then register some further payoffs that this idea has.

First, regarding linguistic data points, consider criticizability in connection with norm violation. When someone violates a norm, we expect criticism to be felicitous or appropriate. It is felicitous and appropriate to say, e.g., “But you didn’t *know* that p!” if it comes to light that someone asserts that p on the basis of a guess, etc.²² That said, suppose A and B purport to be intentionally inquiring together, by means of an aimed conversation, into whether a suspect, Mr. X, is guilty. A and B are amateur sleuths and, while they have a lot of evidence in front of them (say, a box of unsorted documents), they have no good idea where to begin, which questions to start with, what points to make salient. A and B take turns reading scattered bits of the documents out loud, making some arbitrary categorizations of what they’ve learned, asking questions that aren’t pertinent, etc. Carrying on this way, A and B might through this conversation end up by a series of very fortunate breaks reaching the conclusion that Mr. X was innocent (which is correct). But the way they proceeded in conversation is still criticisable here.

²² See, e.g., Williamson (1996) and for an overview Benton (2014).

And note that this would remain so even if, towards the end of their chaotic and disorganized conversation, an expert (Sherlock) gave them reliable testimony that Mr. X was innocent, such that each individually then did *know that* Mr. X was innocent. Even so, they are criticisable on account of the fact that they didn't *know how* to do use conversation so as to take the necessary means to figuring out whether Mr. X was guilty, which is exactly what they purported to be doing, intentionally so. They *used* conversation (their scattershot speech acts met minimal conditions for a communicative exchange), and they reached that end, albeit in a way akin to how a lucky locksmith might pick a lock, the intended outcome, by giggling a pin haphazardly in a way that caused it to open.

The fact that the above inquiry via conversation, absent the relevant participatory know-how, is criticisable offers one line of linguistic data for thinking that participatory know-how is the norm of aimed conversation. Another line of linguistic data is the Moore-paradoxicality of first-person assertions that implicate norm violation. For example, it is not felicitous at all for a jury to say they are going to de-liberate (via conversation) whether someone committed a crime, but then to add that they don't know *how* to use conversation to figure out such a thing – that they don't know what to say to each other or what to ask, or how talking about certain things might help them make any headway.

Setting aside these linguistic points, note that the idea that, e.g., we should use aimed conversation in the service of trying to figure out whether *p* only if we *know* how to use conversation to do this, fits neatly with the very idea implicit in aimed theoretical conversation *that we have a shared intention to figure out whether p*. Recall here the two epistemically relevant points that apply to intentional behaviour generally: that our knowledge of what are doing intentionally will be (by the subject of the relevant actions) characteristically both (i) non-observational and (ii) known by participants to have a *means-end* structure.

If I didn't know how to open a window, and just made movements with my arm near the window and hoped for the best, and it happened to open, and then happened to cast light on the wall, I might learn observationally that I both opened the window and that this resulted in light cast on the wall. But, as per Anscombe, when I open the window intentionally, I know that I open the window in a way that I don't know (as I do merely observationally) that my doing so cast a light on the wall. If we are inquiring into whether *p* by means of a conversation where we satisfy ACK-2 we are *not* like the person who doesn't know how to open the window and learns they did so observationally; and crucially, nor are we like participants in a conversation for conversation's sake and which happens to reach unintentionally the particular denouement whereby those participants learn that Mr. X is innocent. Granted, our conversation might have, for instance, lasted longer than we'd expected, and *that* is something we'd then come to know much like one who intentionally opens the window knows (observationally) that it cast light on the wall. But we would know *non-observationally* that what we are doing is figuring out whether *p*. This point

fits closely with the point about means/end structure: we know we are asking questions, making assertions *in order to* do what we know how to use these assertions and questions to do: to figure out p. Our epistemic position with respect to the connection between our making of these speech acts and our figuring out p isn't the same kind of epistemic position we have simply by observing empirical regularities and inferring an empirical generalisation, e.g., that the making of these kinds of speech acts tends to produce that kind of result, etc. Rather, we know we are using these speech acts as a means to do what we intend to do together (figure out p), viz., we know we are asking and saying what we are asking and saying *in order to* figure out whether p.

In a similar vein, recall the functional point about shared intention. As per Bratman, when you and I have a shared intention that *we* figure out whether p, this shared (participatory) intention is going to structure our bargaining, subplans, etc. Now, just imagine briefly how that bargaining is going if we fail to satisfy ACK-2. Even if you and I are both individually smart enough to evaluate the evidence each of us has at our disposal individually, we still might not know how to structure our conversation together in a way that would facilitate our goal. For an example that brings this point out, consider a conversation between egotistical jurors, ineffective in answering each others' questions, making false assumptions about each others' conversational backgrounds, overemphasizing particular points out of individual ego, not backing down to revisit points, etc. Even if the jurors described here desire and intend to use a conversation to figure out whether the suspect is guilty, *absent know-how* to use speech acts together such as questioning, assertion, etc. in the service of their goal, it's hard to see how their shared participatory intention is going to structure their interactions in a way that will actually lead to their intended *aim*. In this way, satisfying ACK-2 would nicely account for how interlocutors would be in a position to succeed in structuring their bargaining, formation of subplans *in line* with the participatory intention that they share.

Furthermore, consider how the know-how norm under consideration is a natural fit with CAIR-AC*. CAIR-AC* says something about what interlocutors *ought* (the ought of aim-relative, instrumental rationality) to do if intentionally inquiring together whether p by means of an aimed conversation. In particular, what they ought to do is take the necessary conversational means to figuring out whether p. This is something they'll be in a position to do only if they *know-how* to take those necessary conversational means – and so ACK-2 captures this by placing such know-how as a norm on finding out whether p together by means of conversation. In this respect, CAIR-AC* and ACK-2 are complementary normative requirements on aimed conversation, where interlocutors think together to a theoretical purpose.

Finally, consider ACK-2 in connection with individual-level knowledge norms and Gricean norms of cooperation. As the example case involving the Married-Unmarried Riddle suggested, taking the necessary conversational means to figuring out whether p together will involve more (even when individuals have excellent evidence) than just contributing to

the conversation via speech acts that satisfy individual knowledge norms governing those speech acts, along with satisfying Gricean cooperative norms. In such cases, as we've seen, interlocutors might still fail to know how to take the necessary conversational means to figuring out whether p. That said, when ACK-2 *is* satisfied, we should expect that individuals will satisfy the relevant knowledge norms governing their speech acts (e.g., questions and answers) as well as Gricean norms of cooperation. Speech that fails such norms (e.g., false assertions, ill-formed questions) simply won't manifest interlocutors' knowledge how to use a conversation to figure out together the question they are inquiring into.

Putting this together: considerations about instrumental rationality, shared intentionality, the epistemology of intentional action, as well as linguistic data (including concerning criticizability and Moore-paradoxical assertions) all point in the same direction: that aimed conversations are governed by participatory know-how, and that we should intentionally inquire together into whether something is so only if we know how to use conversational means to find out together whether p. These points are reflective of a *know-how* norm on certain kinds of conversations that plausibly entails, but is not entailed by, the satisfaction of more familiar knowledge norms applicable to conversations, and it reveals an important connection between inquiry and a kind of knowledge.

§8.

Might there be such a thing as a knowledge norm governing conversations and not just their constituent parts? Taking this question as a starting point, I began by laying out a precedent found in Williamson for a complex of thought and action to be governed by a knowledge-(implicating) norm. §3-4 argued that a special subset of conversations – (constitutively) aimed conversations – can be norm governed *apart* from the individual speech acts that feature in them. When participants have a shared intention to think together (by means of a conversation) to a particular purpose (e.g., settling a whether p question), the purpose toward which they think together intentionally is partly *constitutive* of the kind of thing they are doing by means of speech acts. §5 then detailed a thesis about instrumental rationality that is applicable to interlocutors engaged in aimed conversation in the service of inquiry: when individuals (A, B, ... n) intentionally inquire together into whether p, by means of a (constitutively) aimed conversation, then (A, B, ... n) ought to take the necessary conversational means to figuring out whether p, and these means will include those structured by the participatory intentions (A, B, ... n) share. §6 then showed that the above requirement on instrumental rationality of interlocutors inquiring by means of an aimed conversation won't be secured *simply* by interlocutors satisfying whatever (series of) knowledge norms govern individual speech acts in the conversation. §7 then defended the idea that aimed conversations are governed by know-how in a way that is not *reducible* to any concatenation of satisfying individual norms on thought and action; the view reached is that individuals (A, B, ... n) must: intentionally inquire together into whether p,

by means of an aimed conversation X, only if A,B,...n *know how* to use X to find out together whether p. I concluded in §7 by taking stock of both the motivations and payoffs of this particular way of thinking about the relationship between knowledge and conversation.

Funding Acknowledgement: Thanks to the Leverhulme Trust (Grant No. RPG-2019-302) and the Arts and Humanities Research Council (Grant No. AH/W005077/1, AH/W008424/1) for supporting this research.

References

- Anscombe, G. 1957. *Intention*. Cambridge, MA : Harvard University Press.
- Beddar, B. and Pavese, C.. 2021. Practical Knowledge without Luminosity. *Mind*, October. <https://doi.org/10.1093/mind/fzab041>.
- Benton, Matthew A. 2014. Knowledge Norms. *Internet Encyclopedia of Philosophy*. <https://iep.utm.edu/kn-norms/>
- Bratman, Michael E. 2013. *Shared Agency: A Planning Theory of Acting Together*. Oxford : Oxford University Press.
- Carter, J. Adam, and Shepherd, J. 2023. Intentional Action and Knowledge-Centred Theories of Control. *Philosophical Studies*, 180 (3):957-977.
- Davidson, D. 1970. *Essays on Actions and Events: Philosophical Essays Volume I*. Oxford: Clarendon Press.
- Friedman, J. 2020. The Epistemic and the Zetetic. *Philosophical Review*, 129 (4), 501–36.
- Gilbert, M. 2008. Two Approaches to Shared Intention: An Essay in the Philosophy of Social Phenomena. *Analyse & Kritik*, 30 (2), 483–514. <https://doi.org/10.1515/auk-2008-0208>
- Gilbert, M. 2009. Shared Intention and Personal Intentions. *Philosophical Studies*, 144, 167–87.
- Grice, H. P. 1975. Logic and Conversation. In: Davidson, D. and Harman, G. (eds.). *The Logic of Grammar*, pp. 64–75. Encino, Calif.: Dickenson Pub. Co.
- Hawthorne, J. and Stanley, J. 2008. Knowledge and Action. *The Journal of Philosophy*, 105 (10), 571–90.
- Ichikawa, J. and Jenkins, C. 2017. On Putting Knowledge “First”. In: Adam Carter, J. Gordon, E. C., and Jarvis, B. (eds.). *Knowledge First: Approaches in Epistemology and Mind*, pp. 113-139. Oxford : Oxford University Press.
- Kearl, T. 2023. What We Know When We Act. *Philosophical Studies*, 180 (9), 2665-2683.
- Kelley, M. 2022. How to Perform a Nonbasic Action. *Noûs*. <https://doi.org/10.1111/nous.12440>.

- Lynch, M. P. 2009. Truth, Value, and Epistemic Expressivism. *Philosophy and Phenomenological Research*, 79 (1), 76–97.
- Mele, A. R., and Moser, P. K. 1994. ‘Intentional Action’. *Nous*, 28 (1), 39–68.
- Miracchi, L. and Adam Carter, J. 2022. Refitting the Mirrors: On Structural Analogies in Epistemology and Action Theory. *Synthese*. <https://doi.org/10.1007/s11229-022-03462-y>.
- Paul, S. K. 2009. How We Know What We’re Doing. *Philosophers’ Imprint*, 9, 1–24.
- Pavese, C. 2015. Practical Senses. *Philosophers’ Imprint* 15.
- Pavese, C. 2021. Knowledge How, April. <https://plato.stanford.edu/entries/knowledge-how/?fbclid=IwAR0Ye8rGs3ImDpf2Vki55VICathCj8El328NYYB534D4EQIwGpbHOIEBDqM>.
- Pavese, C. 2022. Practical Knowledge First. *Synthese*, 200 (5), 375. <https://doi.org/10.1007/s11229-022-03848-y>.
- Piñeros Glasscock, J. S. 2020. Practical Knowledge and Luminosity. *Mind*, 129 (516), 1237–67.
- Schwenkler, J. 2012. Non-Observational Knowledge of Action. *Philosophy Compass*, 7 (10), 731–40. <https://doi.org/10.1111/j.1747-9991.2012.00513.x>.
- Shepherd, J. and Adam Carter, J. [forthcoming]. Knowledge, Practical Knowledge, and Intentional Action. *Ergo: An Open Access Journal of Philosophy*.
- Stanley, J. 2011. *Know How*. Oxford : Oxford University Press.
- Velleman, J. D. 1997. How To Share An Intention. *Philosophy and Phenomenological Research*, 57 (1), 29–50. <https://doi.org/10.2307/2953776>.
- Willard-Kyle, C. 2023. Valuable Ignorance: Delayed Epistemic Gratification. *Philosophical Studies*, 180 (1), 363–84. <https://doi.org/10.1007/s11098-022-01902-6>.
- Williamson, T. 1996. Knowing and Asserting. *The Philosophical Review*, 105 (4), 489–523.
- Williamson, T. 2017. Acting on Knowledge. In: Adam Carter, J. Gordon, E. C., and Jarvis, B. (eds.). *Knowledge First: Approaches in Epistemology and Mind*, pp.163–81. Oxford: Oxford University Press.
- Williamson, T. and Stanley, J. 2001. Knowing How. *The Journal of Philosophy*, 98 (8). <https://ora.ox.ac.uk/objects/uuid:1fcf732b-731b-47e6-943d-aaf34905c02a>.