

Thinking Together

Guy Longworth¹

5.5.22

My affinity, with whom I *think* instead of speak.

—Fanny Seward to, and about, her mother, Frances Seward, February 11, 1864.

1. What is it, or what would it be, for one person to think with another? How should we understand the idea of people thinking together and, in that way or ways, sharing their thinking as well as their thoughts? My primary aim is to begin to explore this question by describing one sort of activity which seems to figure essentially in connection with this more general idea of thinking together.

In our opening quotation, Fanny Seward seems to commit herself to at least two, and perhaps three, relevant claims. The first of these claims is that there is such a thing as *thinking* together. Seward characterizes herself as undertaking this activity of thinking with her mother, Frances. As mentioned, my primary aim will be to uncover one essential element of the sort of activity she may have had in mind. The second claim is that there is such an activity as thinking together with someone *instead of*—that is, presumably, *without*—speaking with them. Although I do not wish to exclude that as a possibility, my present focus will be on the more quotidian seeming cases of thinking together that depend upon speech. Plausibly, however, what Seward really has in mind is not the possibility of thinking together without speaking, but rather the possibility of thinking with someone as opposed to *merely* speaking

¹ I'm grateful for comments or discussion to Anita Avramides, José Luis Bermúdez, Stephen Butterfill, Lucy Campbell, Quassim Cassam, John Collins, Thomas Crowther, Imogen Dickie, Naomi Eilan, Nadine Elzein, Sandy Goldberg, Anil Gomes, Jane Heal, Natalia Waights Hickman, Jennifer Hornsby, Hemdat Lerman, Matthew Parrott, Christopher Peacocke, Ian Rumfitt, Matthew Soteriou, Victor Verdejo, and audiences at Oslo, Oxford, València, and Warwick.

with them—that is, as opposed to speaking with them without, at the same time and thereby, thinking with them. And that possibility then connects with the third of the claims to which Seward seems to commit, that there is a corresponding possibility of merely speaking with someone without thereby thinking with them.

Seward seems not only to commit to the possibility of speaking with someone without thereby thinking with them. In using the fact that she and her mother also think together as a marker of the specialness of their relations, she seems to suggest, in addition, that merely speaking together is a norm of relations with others. By contrast, she seems to suggest, a capacity to think together with someone marks out one's relations with them as distinctively intimate. Once we have a clearer view of the sort of activity that Seward may have had in mind, we will be in a better position to consider the question of whether the activity, or some of its variants, might impose special demands on its participants and the relations amongst them.

I shall begin attempting to reach a clearer view of one important element in some cases of thinking together in the following section by sketching some properties of the putatively wider genus, bare thinking together. The remaining sections aim to sharpen our understanding of the essential element of the sketch that is my primary quarry, the idea of thinking a determinate thought together. That quarry will be pursued via an idea of Ryle's about the ways in which capacities to do things and capacities to appreciate those doings might work together in supporting some sorts of collective activities.

2. Let's begin with the bare idea of φ -ing together, of which the bare ideas of thinking or walking or talking together are instances. To a good first approximation, some individuals are φ -ing together just in case those individuals are the subjects of a measure of *collective* φ -ing. And they are the subjects of a measure of collective φ -ing just in case they are φ -ing and it is not the case that their φ -ing is *distributed* amongst them. Their φ -ing will be collective, and so not distributive, if and only if they are φ -ing and it is not the case that both (i) that they are φ -ing is an immediate consequence of that fact that each of them is φ -ing and (ii) that they are φ -ing is an immediate consequence of the fact that each of them is φ -ing.² Accordingly,

² See, e.g., Oliver and Smiley 2013; Longworth 2019; Butterfill and Sinigaglia 2022. One reason for thinking that this is only an approximation is the appeal to the

Fanny and Frances will be walking together if there is some measure of walking that can be attributed collectively to them, where that requires that they are walking and that at least one of the following claims is true: (i) that each of them is walking does not follow immediately from the fact that they are walking; (ii) that they are walking does not follow immediately from the fact that each of them is. Clearly, that might be so in virtue of various more specific facts about Fanny, Frances, and their interrelations. For example, it might depend on there being measures of walking that are attributable to each of them or it might instead depend only on their each contributing to the collective activity of walking something other than a measure of walking. For an example of the latter sort, they might be walking together because each of them is responsible for only one of the paired steps involved in their walking together. And if their walking together depends on there being measures of walking that are attributable to each of them, the fact that they are walking together might depend on the proximity and interdependence of their measures of individual walking or it might instead depend only on their proximity whilst walking or only on the interdependence of their individual measures of walking. They might be walking together, for example, because they happen to be near to one another while each is walking or because each is contributing a spatially and temporally separate stage in their walking from A to B.³

idea of *immediate* entailment (Oliver and Smiley appeal to analyticity), with the aim of finessing examples wherein, due to the special nature of ϕ -ing, it is either impossible for some number of individuals to ϕ non-collectively without their also ϕ -ing collectively or impossible for some individuals to ϕ collectively without their also ϕ -ing non-collectively. The thought is that where that is so, the target entailments are derivative rather than immediate, running as they do via the more immediate connections between collective and non-collective ϕ -ing. An additional reason for thinking that this is only an approximation concerns the relations between togetherness and collectivity. As Annette Baier suggests (1997: 28), ϕ -ing together, by contrast with bare collective ϕ -ing, often seems to require co-operation or, at least, the absence of non-cooperation. Crudely, ϕ -ing together requires having a collective end, which may just be ϕ -ing, or at least absence of immediately conflicting ends. For present purposes, the relevant claims about the relations between togetherness and collectivity can be treated as stipulations about how “ ϕ -ing together” is to be understood here.

³ Whether bare proximity suffices for togetherness might be disputed or might be treated as a special case. One thought is that in some cases, we allow that A and B are ϕ -ing together only because A and B are ψ -ing together while each of them is separately ϕ -ing. So, for example, we might allow that A and B are walking together only because they are talking together while each of them is separately walking or because their walking separately serves their collective ends. In a similar vein, it might be that we sometimes allow that A and B are

Similarly, Fanny and Frances might be thinking together because they are together while each of them is thinking, where that seems to be a possibility even if each of them is thinking silently.⁴ Alternatively, they might be thinking together because each of them is contributing a separate stretch of thinking to a more extended measure of thinking—for example, by their separately contributing the derivations of distinct lemmas to a more expansive derivation or by one of them asking questions to which the other supplies answers.

What about the thinking of lemmas? That is, what about acts of thinking determinate thoughts—for example, acts of thinking the thought that Fanny is bored? Are acts of that sort examples of things that Fanny and Frances can do together? And if collective acts of thinking determinate thoughts are possible, how should we understand that possibility? Alternatively, if collective acts of thinking determinate thoughts are *not* possible, then we might worry that that would throw into question the possibility of thinking together more generally. For we might hold that for a stretch of thinking to be genuinely collective, there must be points in that thinking that are themselves genuinely collective. More modestly, we might hold that genuinely collective thinking must exhibit at least the potential for collective thinking of determinate thoughts, and so is dependent on the possibility of point-wise collective thinking. Otherwise, although we could still allow that the individual thinking of one of them had bearing on the individual thinking of the other, we might hesitate to allow that they were really thinking together as opposed, say, to doing only other things together while each of them was thinking separately. I shall assume in what follows that this more modest position is correct, so that thinking together is dependent on the possibility of acts of thinking determinate thoughts together. As in the case of thinking and walking, we might want to allow that Fanny and Frances can together be thinking that Fanny is bored because they are together in some other way while each of them thinks separately that Fanny is bored. But that would not be a case in which there was a single act of thinking that Fanny is bored that was undertaken jointly by Fanny and Frances. Could there be a case of the latter sort?

Given the discussion to this point, there are two natural paths along which to seek for a case in which Fanny and Frances jointly

walking together only because A and B are spatially together (barely proximal) while each of them is separately walking.

⁴ Here, again, it might be argued that we allow this only because they are, say, walking together while each of them is separately thinking.

undertake a single act of thinking that Fanny is bored. The first path is marked out by each of Fanny and Frances contributing to their joint activity an individual act of thinking that Fanny is bored. The second is marked out by each of them contributing something other than an individual act of thinking that very thought. Let us start by considering the prospects of the second path.

The second path branches again. Along the first of its branches, each of Fanny and Frances contributes an individual act of thinking a thought, but it is a thought other than the thought that Fanny is bored. The main obstacle to our following that branch is that it is difficult to see how each of Fanny's and Frances's thinking thoughts other than the thought that Fanny is bored could eventuate in their jointly thinking that very thought. For example, we might suppose that Fanny contributes by thinking the thought that Fanny is some way whilst Frances contributes by thinking the thought that someone is bored. But it is hard to see how their thinking jointly that Fanny is bored can comprise those two individual stretches of thinking. What seems to be required is the thinking of a thought that brings together what Fanny and Frances each think. And that seems to require, in turn, that there be some individual thinking of the required thought.

It might be supposed that this issue could be finessed by appeal to their individually thinking thoughts which, in combination, manifestly entail that Fanny is bored. The idea would be that in the case of an individual's thinking such thoughts, we might be prepared to allow that their doing so can support attributing to them the thought that Fanny is bored. Similarly, then, the idea would be that where Fanny and Frances each think thoughts that together entail that Fanny is bored, even though neither Fanny's thinking nor Frances's carries that entailment on its own, we might similarly be prepared to allow that their doing so can support attributing to them collectively the thought that Fanny is bored.

The attempt to finesse the issue in this way is subject to numerous difficulties, of which three seem most pressing. The first of these difficulties is that the precise analogy that the attempt seeks to work between individual and collective cases seems to depend on the same subject or subjects being responsible for thinking both the entailing thoughts and the entailed thought. It is plausible that an individual thinks the entailed thought only because that very individual thinks the entailing thoughts. And yet the attempt fails manifestly to meet that condition, since it involves only each individual thinking separately the entailing thoughts and only the individuals together thinking the entailed thought.

The second difficulty is related. One natural explanation for why it can be reasonable to take subjects who think some entailing thoughts to also think an obviously entailed thought is that such subjects will normally be *disposed* to think the entailed thought. Since being so disposed is ordinarily a basis for being held to think the thought in dispositional form, it is natural to allow that where someone would normally be so disposed, they think the entailed thought. However, this is another point at which the analogy between individual and collective cases breaks down. Even if we ignore the first difficulty, and so focus on occurrent thinking of an entailed thought by any of the individual thinkers, the transition from distributed thinking of the entailing thoughts to individual thinking of an entailed thought requires to be mediated by more than normality. Plausibly, it requires, in addition, is that the thinking of the entailed thoughts be brought about by the thinking of the entailing thoughts. And that seems to require communication of all the entailing thoughts to at least one thinker.

The third difficulty facing the attempt arises when we recombine the first and second difficulties. Minimally, what the attempt is required to sustain is the possible collective possession of a disposition to *occurrently* think an entailed thought. However, that possibility can be sustained only if it is possible for subjects to occurrently think a thought collectively. The attempt simply presupposes that possibility, and so goes no distance at all towards underwriting it. Plausibly, occurrent thinking essentially involves occurrent vehicles of thinking, and so making out the possibility of collective examples of occurrent thinking would require explaining how there can be collective exploitation of occurrent vehicles of that thinking.⁵ We will have cause to return to the question of whether, or how, that is possible.

We are still considering the path to thinking a determinate thought together along which participants contribute something other than an act of thinking that very thought. The first of that path's branches, according to which participants nonetheless contribute an act of thinking a thought, seems to be blocked. What about the second branch, on which participants contribute something other than such an act? Here, the most natural proposal to consider is one on which participants collectively contribute analogues of whatever underpins individual acts of thinking. Thus, suppose that where Fanny thinks that Fanny is bored, her thinking is underpinned by brain state B_1 and brain state B_2 , as appropriately connected in Fanny's overall neural economy. The proposal is that

⁵ On occurrent thinking more generally, see Soteriou 2013

where Fanny and Frances think together that Fanny is bored, their thinking is underpinned by Fanny's being in brain state B_1 and Frances's being in brain state B_2 again, as appropriately connected in a system comprising at least Fanny's and Frances's respective neural economies.

One obvious difficulty with this proposal is the difficulty of understanding how Fanny's and Frances's respective brain states can be appropriately connected—connected, that is, in a way appropriately analogous to the way in which Fanny's own brain states are connected. However, an even more pressing difficulty is that insofar as Fanny's and Frances's brain states together simulate the states of Fanny's brain when she thinks that Fanny is bored, and Fanny and Frances together thereby contribute to the simulation of Fanny's thinking that Fanny is bored, they contribute to the simulation of an *individual's* separately thinking that thought. For it is an individual's thinking that the collective simulates.

A related problem arises from the opposite direction. The act of thinking that Fanny is bored—and, more generally, the act of thinking a determinate thought—seems to demand at least one, unified subject. That is, for some people collectively to instance a measure of thinking, at least one of their number, and perhaps each of their number, must contribute by separately instancing a measure of thinking. Because of this demand, there is an additional reason to understand Fanny and Frances in the target case as contributing not to a collective act of thinking that Fanny is bored, but rather to a collective act of constituting a single, unified subject's thinking that Fanny is bored. For otherwise, since neither Fanny nor Frances engage in the target act of thinking, that act would lack a unified subject.⁶ If that is right, then it plausibly follows from the fact that Fanny and Frances are thinking together that Fanny is bored that at least one of Fanny and Frances, and perhaps each of them, must be thinking that Fanny is bored.⁷ What remains as a potential marker of the collectiveness of their thinking is just that their thinking together that Fanny is bored doesn't follow merely from that fact that each of them thinks separately that Fanny is bored.

The second path seems to be closed. If it is possible for Fanny and Frances to think together that Fanny is bored, their doing so

⁶ At this stage, we might find ourselves wanting to say something similar about the act of walking and, on that basis, wanting to revisit our earlier judgement that Fanny and Frances could walk together by each contributing less than a measure of walking to their collective endeavour.

⁷ That conclusion incidentally bolsters the obstacles attending the other branch of this path.

must involve at least one of their thinking that Fanny is bored. The difficulty that attends explaining the possibility of their thinking together that Fanny is bored is therefore the difficulty of explaining what their doing so can involve over and above their separately thinking that thought. What is wanted is an account how Fanny's thinking can connect appropriately with Frances's, through the interlocking exercises of their respective abilities to think.

3. The previous section refined our quarry. In seeking a case of thinking together, part of what we seek is a collectively instanced measure of thinking a determinate thought. Making out the possibility of such an instance depends on making out the possibility of a collective instance of the occurrent thinking of a determinate thought. And that in turn depends, first, on explaining what the collective thinking of a determinate thought involves over and above the separate thinking on which it depends and, second on making out the possibility of there being an occurrent vehicle for that collective thinking. We shall now begin cautiously to approach our quarry through developing some aspects of a discussion by Jennifer Hornsby (2005) of our abilities to voice our thoughts, which builds, in turn, on an earlier discussion by Gilbert Ryle (1949).

The aspects of Hornsby's discussion on which we will focus develop an underappreciated observation of Ryle's, that there is a sort of symmetry between knowing how to do something and being able, or having the cognitive wherewithal, to tell whether someone else is doing that thing (see Hornsby 2005: 124-5). Ryle's observation will figure importantly in explaining how individual's abilities to think can interlock so as to sustain their thinking together.

Understanding what someone is up to, Ryle suggests, does not depend on risky inference or guesswork about their inner life, and that suggestion raises a question:

If understanding does not consist in inferring, or guessing, the alleged inner-life precursors of overt actions, what is it? If it does not require mastery of psychological theory together with the ability to apply it, what knowledge does it require? (Ryle 1949: 53)

Ryle's answer to this question exploits his symmetry observation:

We saw that a spectator who cannot play chess also cannot follow the play of others; a person who cannot read or speak Swedish cannot understand what is spoken or written in Swedish; and a

person whose reasoning powers are weak is bad at following and retaining the arguments of others. Understanding is a part of knowing *how*. The knowledge that is required for understanding intelligent performances of a specific kind is some degree of competence in performances of that kind. (Ryle 1949: 53)

Ryle's suggestion here is that one's understanding another's intelligent performances—that is, one's understanding another's exercises of knowledge how to do something—depends upon one's possession and exercise of the same form of knowledge how. Being able to understand spoken Swedish, for example, depends upon knowing how to speak Swedish.

Further, and in partial explanation of the suggestion, Ryle claims that the ability to understand acts of some kind is a *part of* knowing how to perform acts of that kind (or a part of being able to perform them, at least intentionally). The core claim here is that knowing how to ϕ is itself dependent on being able to understand acts of ϕ -ing as such and, in that sense, the latter ability is part of knowing how to ϕ . And that claim is plausible, insofar as one's knowing how to ϕ plausibly depends both on one's knowing what it is to ϕ and on one's being able to tell whether one is ϕ -ing or has ϕ -ed. To a good first approximation, the idea is that knowing how to do something requires one to know what it is to do that thing, and to possess that knowledge in a form that connects appropriately with the concrete vehicles of performing acts of that sort—e.g., bodily movements. But such knowledge plausibly provides its possessor with the cognitive wherewithal not only to trace a path from intention to concrete implementation but also to follow that path back from concrete implementation to intention.

However, the core claim, that the ability to understand is a part of the ability to act, admits of more than one development. On one way of developing it, the idea would be that the ability to understand is a *detachable* component of the ability to act. That is, one cannot know how to ϕ without being able to recognise ϕ -ing, but, since the latter ability is detachable, one can be able to recognise ϕ -ing without knowing how to ϕ . That way of developing Ryle's core claim is inadequate to his own larger purpose. It would sustain the result that knowing how to ϕ *suffices* for being able to understand ϕ -ing, since the latter would be a necessary component of the former. But Ryle's larger purpose is to support the suggestion that knowing how to ϕ is *necessary* for being able to understand ϕ -ing and so achieving that purpose depends upon his rejecting the idea that the ability to understand ϕ -ing can be possessed by one

who does not know how to ϕ . What Ryle needs, then, is a way of developing the core claim which delivers the result that the ability to understand is a non-detachable component of knowing how.

Whilst it is comparatively easy, as we have noted, to see, in outline, why knowing how to ϕ should depend on being able to recognise cases of ϕ -ing, it is less easy to see why the latter ability should depend on the former knowledge. Isn't it possible, say, to be able to recognise that another is skiing without oneself knowing how to ski?⁸

The issues here are delicate, and we should avoid leaping to the conclusion that Ryle's necessity claim is indefensible. For one sort of approach, we might consider following Michael Dummett in suggesting that a version of Ryle's necessity claim is true of a restricted class of activities. For example, we might follow Dummett in allowing that

knowledge of a language...is an acquired ability to engage in a practice of such a kind that one cannot know what engaging in it consists in until one has acquired the ability to do so. (Dummett 1991: 94)

Dummett does not have much to say about the special features of the practice of speaking a language that mark it out for this seemingly special treatment. One thought here might be that the operative feature of the practice of speaking a language is that one cannot engage in it unintentionally, since the practice is partly constituted by a distinctive range of intentions with which participants engage in it. However, even if that suggestion were defensible, we would still want an account of why that feature dictates that one cannot know what engaging in the activity consists in without having acquired the ability—or better, the know how—so to engage.

⁸ It is harder still to see why the ability to recognise cases of ϕ -ing should depend on an *ability* to ϕ . Couldn't one, for example, preserve one's ability to recognise that someone is skiing having lost one's own ability to ski? Well, that might depend upon precisely which abilities are at issue: we often allow that someone can preserve an ability to ski through temporary disability. One suggestion would be that what is disabled, in some such cases, is only what is physically required (bodily or environmentally) for skiing here and now, while what is retained is the ability intentionally to ski—that is, the cognitive wherewithal to ski—insofar as various physical enabling conditions are in place, where having the latter ability is necessary for being able to recognise others as skiing. See, e.g., Hornsby 2005: 115, including fn.15.

Happily, however, I can afford to remain neutral here about the resolution of the delicate issues that surround Ryle's necessity claim, for my purposes are not the same as his. My purposes depend only on defending a form of Ryle's sufficiency claim, according to which knowing how to φ requires, or includes, an ability to recognise someone's φ -ing.

We saw that Ryle's sufficiency claim is plausible, at least in outline. However, it might be doubted on the following grounds. An ability to recognise someone's φ -ing is an ability to know, of someone who is φ -ing, that they are φ -ing. When this knowledge is knowledge of someone else's φ -ing, the form it takes seems to be distinctively *theoretical* or *receptive*: the knowledge is determined as such from without—through, for example, observation of someone's φ -ing. By contrast, when this knowledge is knowledge of one's own intentional φ -ing, the form it takes seems to be distinctively *practical* or *originative*: the knowledge determines itself as such by bringing about its object, the fact that one is φ -ing. Furthermore, knowing how to φ seems, in the first instance, to sustain only knowledge of one's own (intentional) φ -ing. And so, if this difference in form between, on one side, practical knowledge of one's own φ -ing that can be sustained by one's knowing how to φ and, on the other side, theoretical knowledge of another's φ -ing makes for a difference in what one knows in each case, then there would be reason to doubt that knowing how to φ is sufficient for being able to recognise that another is φ -ing.

The thought that that is so—that what one knows when one has practical knowledge that one is φ -ing is distinct from what one knows when one has theoretical knowledge that someone is φ -ing—might be supported by the thought that practical knowledge embodies a distinctively practical perspective on one's actions. In that case, although knowing how to φ would enable one to know whether one is φ -ing, there would be no immediate reason to expect it to enable one to know whether someone else is φ -ing. For the knowledge that one was able to have of one's own actions would embody a perspective on those actions that would make one's ability non-transferrable. It could not sustain the kind of perspective on one's own actions that would be required for one to thereby be positioned to enjoy theoretical knowledge of those actions, and so it would not put one in a position to recognise another's actions of the same kind.

Although tempting, that line of thought is mistaken. To do the work required of it by the argument sketch, the idea that practical

knowledge embodies distinctively practical perspectives on one's own actions would need to sustain the conclusion that the relations between pieces of practical knowledge and pieces of theoretical knowledge are opaque. Differences in the perspectives embodied by one's practical knowledge that one is ϕ -ing and any theoretical knowledge that one is ϕ -ing would need to play the same sort of role in blocking transparent connections amongst pieces of knowledge as is played by differences in the perspectives embodied by knowledge that Hesperus twinkles and knowledge that Phosphorus shines. In the latter case, one possessing both of those pieces of knowledge would be unable, just on their bases, to conclude knowledgeably that something twinkles and shines. Similarly, one with theoretical knowledge that if they are bowing, then they win, and practical knowledge that they are bowing, would not thereby be positioned to conclude knowledgeably that they win. Relatedly, it is plausible that one who knows only that Hesperus twinkles cannot enable an audience to know that *Phosphorus* twinkles by telling their audience either that Hesperus twinkles or that Phosphorus does. And similarly, it is plausible that one with practical knowledge that they are leaving, or soon will be, would not be able to let another know that they are by telling them, for their being able to let the audience know that they are through knowledgeable testimony to that effect would seem to depend on their sharing with their audience a perspective on their leaving.⁹ Since we are not in fact subject to the various disabilities predicted by the claim that practical knowledge contrasts with theoretical knowledge in embodying only a distinctively practical perspective, that claim should be rejected. The differences between practical and theoretical knowledge provide no grounds for doubting Ryle's sufficiency claim.¹⁰

We have considered two forms of symmetry claim concerning the relations between knowing how to ϕ and being able to tell whether someone is ϕ -ing. Both embody the claim that there is a form of being able to tell whether someone is ϕ -ing that figures in the same way in knowing how to ϕ and in being able to tell whether someone else is ϕ -ing. Ryle's strong form of the symmetry claim involves his necessity claim, that the *only* way of being able to tell

⁹ For relevant discussion, see Dummett 1975; Heck 1995.

¹⁰ Compare Anscombe:

If I say I am going for a walk, someone else may know that this is not going to happen. It would be absurd to say that *what* he knew was not going to happen was not the very same thing that I was saying *was* going to happen. (1957: §52.)

whether someone else is φ -ing depends on knowing how to φ . We have adopted only the weaker form of the symmetry claim. According to this weaker form, *one* way of being able to tell whether someone else is φ -ing is through knowing how to φ . Knowing how to φ involves, at least in part, knowing what it is concretely to φ , where the perspective embodied in that knowledge makes it applicable both to one's own intentional φ -ing and to others' φ -ing.

Ryle acknowledges the complementary roles that knowing how to φ plays, in φ -ing and in engaging with another's φ -ing, in the following passage:

Of course, to execute an operation intelligently is not exactly the same thing as to follow its execution intelligently. The agent is originating, the spectator is only contemplating...He is merely thinking what the author is doing along the same lines as those on which the author is thinking what he is doing, save that the spectator is finding what the author is inventing. The author is leading and the spectator is following, but their path is the same. (Ryle 1949: 53-54)

We see the idea here of there being a distinction between, on one hand, active or originative exercises of know-how and, on the other hand, passive or receptive exercises of that knowledge. Ryle's thought is that the differences between active and passive here are differences in the way a single piece of knowledge is used, and do not reflect the exercises of distinct pieces of knowledge. In the following three sections, we will consider how the interlocking, active and passive exercises of shared know how might sustain the collective thinking of determinate thoughts. In the following section, we will consider receptive exercises of know how before turning, in §5, to active exercises and, in §6, to their combination.

4. We often characterize ourselves and others as hearing (or seeing or feeling) people say things. For example, I might say that I heard Fanny say that she was bored, and you might agree that you heard her say it too. These characterisations are akin to other cases in which untensed small clauses, like "Frances cough," are used to denote events, as in "I heard Frances cough." To a good first approximation, the last example can be represented as in (1):

- (1) There was an event e of coughing, with Frances as subject, and I heard e .¹¹

Similarly, it would be natural to represent (2) as in (3):

- (2) I heard Fanny say that she was bored.
- (3) There was an event e of saying that she was bored, with Fanny as subject, and I heard e .

If that is right, then hearing someone say such and such is a form of hearing an event. And, crucially, hearing an event is standing in an extensional, perceptual relation to a concrete occurrence. So, for example, if the event of Fanny's saying that she was bored was the event of her saying the last thing that she said before bed, then (2) and (4) are equivalent:

- (4) I heard Fanny say the last thing that she said before bed.

Plausibly, however, being the last thing that Fanny said before bed is not a fundamental kind of event: there being an event of that kind depends upon there being a more basic answer to the question, What happened? Plausibly, in this case, the most basic answer to that question is that Fanny said that she was bored. That plausibly represents the fundamental kind of the event one heard.

Now hearing an event of some specific, fundamental kind, for example an event in which someone says something, depends obviously on the possession and exercise of a general ability to hear. It depends, in addition, on whatever more specific auditory and other sensitivities are required to hear events of this fundamental kind in this specific type of circumstances. One's being able to stand in extensional, perceptual relations to concrete occurrences of a specific, fundamental kind depends on one's being appropriately sensorily receptive to occurrences of that kind. However, one can hear an event of a specific kind without recognising it, and certainly without recognising it as an event of that specific kind. That is, one can have heard an event of some fundamental kind without knowing that it is of that kind. One can have heard Fanny say that she was bored, for example, without recognizing that she said that she was bored. The ability to hear events of a fundamental kind is distinct from the ability to recognize events as of that kind.¹²

¹¹ See, especially, Higginbotham 1983.

¹² We can leave open whether or not these abilities are related, perhaps through the ability to hear events of a kind being part of the recognitional capacity or

Recognizing that Frances coughed, on the basis of hearing Frances cough, can seemingly involve the immediate application of a recognitional ability. That is to say, having heard the event of Frances' coughing, it is possible for someone with an appropriate recognitional ability to recognise that Frances coughed without first needing to recognize other features of the heard event and then base their view that Frances coughed on their knowledge of the presence of those other features. We can put this thought by saying that recognizing that Frances coughed through hearing her cough is something that we can *simply* do.¹³

Similarly, we should find it plausible that recognizing that Fanny said that she was bored, having heard her say it, is something that we can simply do. That is, we should find it plausible that if we are to recognize that Fanny said that she was bored, then we do not first need to recognise that the event of her saying that she was bored has some other features—say, specific auditory or grammatical features—and then infer from the presence of those features that Fanny said that she was bored. If we are to avoid regresses, then some such simple recognitional achievements must be possible. It must be possible for us to recognize the obtaining of some ways for things to be without our doing so depending, inferentially, on our prior recognition of the obtaining of other ways for things to be. Given that there must be such abilities simply to recognise the obtaining of some ways for things to be, it is not clear why these abilities should not include the ability simply to recognize what someone has said.

In considering the idea that those of us with the right recognitional capacities can simply recognise what someone has said, it is important to separate that idea from two potential accretions. The first of these accretions holds that it would follow from our being able simply to recognise what someone has said that we could do so without needing to hear, or otherwise perceive, the event of their saying it. As we have already noted, it is possible to hear the event of someone's saying something without recognizing that it is an event of that kind and, plausibly, without there being some specific range of features of the event that one must have recognized. The idea that we can simply recognize what someone said requires only that we can do so without our doing so depending on inference from any other features of the event that we have

through both hearing and recognition being sub-capacities of a more general capacity that incorporates both hearing and recognition.

¹³ Here, I extend Hornsby's analogous use of "simply" (e.g., 2005: 114) which will be discussed further below.

recognized. Since one can plausibly hear the event without one's doing so involving any specific range of exercises of recognitional capacities and, on that basis, exercise immediately some of one's recognitional capacities, the idea of simple recognition is compatible with its being dependent on the perception of events. Further, and more positively, the idea of a simple recognitional ability is the idea of an ability that is exercised immediately given specific forms or types of perceptual experiences, rather than being exercised only mediately, on the bases of prior exercises of recognitional abilities. The idea that there are such abilities is not, therefore, the idea that they can be exercised independently of any other abilities, including perceptual abilities, and it is not the idea that their exercises are independent of ongoing perceptual opportunities. Rather, it is the idea that such abilities can be exercised on bases other than the products of exercises of abilities of the same, recognitional kind.

The second accretion involves the idea that if we have an ability simply to recognize what someone says, then it must be that we either do not, or cannot, recognize other features of events of someone's saying something, and certainly that we cannot recognize such features of events that would enable us to attain mediated knowledge of what they said. What matters to the idea of simple recognition of what someone says is just that the successful exercise of such abilities does not depend on inference from the outputs of exercises of other recognitional abilities. That leaves it open that one who exercises such an ability will, or even must, at the same time exercise other recognitional abilities. And it also leaves it open that it would be possible, either in principle or for a particular subject here and now, to figure out what someone says from knowledge of other features of a target event that was delivered by other recognitional abilities. That someone knows what someone says and that it would have been possible for them to know it by exercising some proper sub-set of their abilities does not entail that they know it on bases supplied by the exercises those abilities. It leaves open that they know it on bases supplied by the exercises of other abilities, including abilities simply to recognize what someone says. Furthermore, the claim that there are abilities simply to recognize what someone has said is compatible with the view that successful exercises of those abilities are dependent upon sensitivities to collateral circumstances, facts, or reasons, even where those sensitivities take the form of exercises of further recognitional abilities. What matters to the idea of abilities simply to recognize what someone says is that any such sensitivities figure

only as enabling conditions for the successful operation of such abilities, rather than as supplying bases for their inferential operation.¹⁴

Having separated the idea of there being abilities simply to recognize what someone says from some potential accretions, I think it is plausible that the idea represents a possibility. It is possible that some people have abilities simply to recognize what someone says given only that they heard them say it. However, it is worth noting that that possibility might take different forms.

First, we might want to allow that it is possible barely to perceive an event without yet perceiving enough of the event, or enough of its concrete features, to be able successfully to exercise relevant recognitional abilities. Perhaps, for example, it is possible barely to hear Fanny say that she is bored without sensorily registering features of that event on which the successful operation of one's ability to recognize what she said depends. So, we might want to allow that there is room for further strictly perceptual abilities to intervene between bare perception of an event of someone's saying something and what is needed successfully to exercise an ability simply to recognize what they said.

Second, there might be kinds of abilities the successful exercises of which was required for simply recognizing what someone said but which are neither strictly perceptual abilities nor full-blown recognitional or knowledge-delivering abilities. For example, I have suggested in other work that a capacity to be induced by the audible expression of thoughts to undergo a receptive mode of thinking those thoughts might sometimes figure between on one side, merely hearing someone say something and, on the other side, recognizing what someone has said (Longworth 2018). Just as the proper operation of an ability simply to recognize what someone says might be sensitive to the outcomes of exercises of perceptual abilities beyond those needed simply to perceive an event of their saying it, so it might be sensitive to the outcomes of exercises of non-perceptual abilities. All that is excluded is that those abilities are themselves recognitional abilities that provide inferential bases for the operation of the ability to recognize what someone says.¹⁵

¹⁴ For relevant discussion, see McDowell 1994.

¹⁵ Having allowed the possibility that some recognitional capacities might depend on the latter sort of capacity, we might also consider whether that sort of capacity figures essentially in the *bare* perception of events in which people say things. Perhaps, for example, all that can barely be perceived without exercising such capacities is more minimal events—e.g., in which sounds are produced.

I now want to begin developing a line of thought in support of a specific version of this latter sort of view as an alternative to the idea that recognition operates immediately only on the outcomes of exercises of strictly perceptual abilities. I shall do so via discussing some ways in which people can voice their thoughts and by bringing to bear on that discussion the symmetry claim that was developed in the previous section. I shall begin by returning to some further relevant aspects of Hornsby's discussion of knowledge of how to voice one's thoughts.

5. As I noted earlier, Hornsby follows Ryle in emphasising certain apparent symmetries amongst our abilities to speak and to understand others' speech.

Just as it seems that you directly hear utterances as the meaningful things that they are, so it seems that when you speak you directly produce meaningful things: it seems...that you directly *voice your thoughts*. (Hornsby 2005: 112)

In the previous section, we provided a way of understanding the claim that some of us are able directly to hear utterances as the meaningful things that they are. On that way of understanding the claim, it requires disambiguation. It is certainly possible for us to hear meaningful utterances. However, that is a matter of hearing events and, as such, a matter of standing in perceptual relations to events together, perhaps, with some of their concrete features. However, I take it that what Hornsby has in mind by hearing utterances as the meaningful things that they are is not a matter only of standing in perceptual relations to events of utterances and some of their concrete features. Rather, what Hornsby has in mind is hearing of a particular utterance *that* it means what it does. That is, what she has in view is the seeming possibility of knowledge that is based on (and perhaps partly constituted by) one's hearing of events. So, when Hornsby says that it seems that one directly hears utterances as the meaningful things they are, what she has in view is that one seems to recognize what utterances mean, or what people say in producing those utterances, without this depending on inference from other facts that one has recognized to obtain. Symmetrically, she suggests, something closely similar holds of our abilities to voice our thoughts.

Hornsby articulates her view about (some of) our abilities to voice our thoughts in the following way:

Speakers can rely on the fact that producing meaningful things is something that they are able to simply do. When a speaker says that *p*, there need be nothing such that she intentionally does it and says that *p* by doing it. (Hornsby 2005: 118)

Hornsby's view is motivated, in part, by the idea that there is the same sort of need to avoid regresses in the case of acting as there is in the case of recognizing. Suppose that, for every act, it were possible intentionally to perform that act only by intentionally performing some other act. In that case, we would face a regress and it would be impossible ever to get started by intentionally performing any act. If it is possible for us to do anything intentionally, as it obviously is, then it must be possible for us to do some things intentionally without doing so by doing other things intentionally. In that sense, there must be some things that we are able *simply* to do (Hornsby 2005: 114).¹⁶ Since there must be some things that we are able simply to do, we should be open to the possibility that voicing the thought that *p* by saying that *p* is amongst those things: some of us have the ability simply to voice our thoughts.

An important, more specific reason that Hornsby offers for thinking that voicing our thoughts is amongst the things we can simply do derives from the connection that obtains between one's acting intentionally and one's having non-observational knowledge of what one is doing. As famously emphasised by Anscombe, one mark of intentional action is that one who acts intentionally thereby knows without (further) observation what they are doing.¹⁷ It follows that where one does not know without observation something that one is doing, that is a reason for thinking that one is not doing that thing intentionally.

Hornsby applies this connection between intentional action and non-observational knowledge to the case of voicing one's thoughts in the following way. If voicing one's thoughts were something that one could do only by intentionally doing something

¹⁶ For relevant discussion, see also Small 2019.

¹⁷ Anscombe 1957. Observation may play a sort of enabling role in one's being able to act intentionally, and so in one's knowing what one is doing when one so acts. In line with our discussion in the previous section, what is excluded here is that one's knowledge of what one is doing is based either solely on observation or on inference from what one observes. Amongst the possibilities that are excluded is that one might act intentionally without knowing what one is doing simply because one failed, for whatever reason, to observe what one was doing. See, e.g., Anscombe 1957: §28.

else, then in those cases in which one voices one's thoughts, the expectation would be that one would know without observation not only that one was voicing one's thoughts but also that one was doing whatever else one needed to do to voice one's thoughts. And yet it seems obvious that one can voice one's thoughts without knowing very much at all about whatever one is doing other than voicing one's thoughts. It seems that one can voice one's thoughts without having non-observational knowledge of anything specific that one has done in order to do that—without, that is, having non-observational knowledge of one's producing specific sounds, or vocal gestures, that one might take to be means of voicing the specific thoughts that one does. Fanny can say that she is bored without having non-observational knowledge of whatever specific acts she uses her vocal system to perform in order to say that. She can know without observation that she voiced the thought that she was bored without knowing in the same way whatever else she did to voice that thought. That pattern of non-observational knowledge and its absence signals that Fanny voiced the thought that she was bored without there being anything else that she did such that she voiced her thought by intentionally doing that other thing. Fanny simply voiced her thought. (Hornsby 2005: 121–123.)

I agree with Hornsby that we can simply voice our thoughts, that this is something we are able to do. However, like the idea that we are able simply to recognize what someone says, the idea that we are able simply to voice our thoughts can attract accretions. We can separate the idea from some of these accretions by considering a putative objection that has been pressed by Jason Stanley (2005) and Richard Kimberly Heck (2006).

The objection, as developed by Stanley, takes off from the observation that awareness of the perceptible features of speech that one exploits in voicing one's thoughts can figure importantly in much of what one does in speaking. Such awareness can figure, for example, in one's planning how to say what one wants to say, or in one's modulating the specific ways in which one says it. Where one's awareness figures in this way in one's voicing one's thoughts, it might seem to figure in delineating means that one intentionally exploits to voice the specific thoughts one does. It anyway seems plausible that one sometimes intentionally exploits such features in this way as means to voicing one's thoughts. And we can add that where one does so, one plausibly does have non-observational knowledge of what one is doing, as required for this exploitation of means to be amongst the things one does intentionally. And so, insofar as one makes intentional use of the perceptible features of

speech as means to one's saying what one does in the way that one does, it is not true that one simply voices one's thoughts. For in such cases, there are things other than voicing one's thoughts that one does intentionally in order to voice one's thoughts.

Suppose that we were to take at face-value Stanley's observation that we often do exploit a sensitivity to features of what we do to voice our thoughts, and that we do so both in determining the thoughts we voice and the ways in which we voice them. Still, for that observation to serve as the basis of an objection to the possibility of simply voicing one's thoughts, further conditions would need to be met. First, and most pressing, it would need to be shown not only that we often do exploit such sensitivities in voicing our thoughts, but also that we *must* exploit them. For the claim that we can simply voice our thoughts is not, on its face, the claim that *whenever* we voice our thoughts, we do so without intentionally exploiting means to do so. In discussing the possibility of simply recognising what someone says, we noted the need to separate it from an accretion, according to which for that to be possible, it must be *impossible* for recognition to be mediated. A similar accretion might seem to be operative here, in the idea that for it to be possible simply to voice thoughts, it must be *impossible* to voice thoughts by intentionally doing something else. But there is no reason to think that one cannot have an ability simply to do something at the same time as having an ability to do that thing by doing something else, and so Stanley's observation sponsors no immediate objection to the possibility of simply voicing thoughts.

Although that suffices to insulate Hornsby's claim from Stanley's observation, it is worthwhile to note a further gap that would need to be traversed to connect observation with claim. This gap falls between the observation that in voicing our thoughts we exploit a sensitivity to ways and means and the further claim that we intentionally exploit ways and means. For we can allow that our awareness of ways and means figures in enabling or partly constituting our intentional activity without being forced to accept that it reflects distinct intentions directed upon those ways and means. Importantly, that can be so even where our awareness of ways and means takes the form of non-observational knowledge. We might compare here the way in which our non-observational knowledge of limb disposition plays an enabling or partly constitutive role in our abilities intentionally to reconfigure that disposition.¹⁸

¹⁸ For relevant discussion, see McDowell 2011; Small 2019.

Hornsby's claim, that we are able simply to voice our thoughts, withstands exposure to Stanley's observation. However, her claim can be understood in different ways, depending on how, exactly, we understand the core idea of voicing one's thoughts. In the following section, I shall distinguish two such understandings and consider how each of them interacts with Ryle's weaker symmetry claim.

6. The first way of understanding the idea of voicing one's thoughts focuses on the activity of *publicising* one's thinking. Cases of this sort would be cases in which one has undertaken a piece of thinking independently of giving voice to it. Since we are focusing on occurrent thinking, this thinking will require a conscious, occurrent vehicle, perhaps a train of silent soliloquy, or other pertinent imagining. One more specific sort of case would involve asking oneself a question, perhaps concerning Fanny's state of mind, with the aim of bringing an answer to mind. In response, one might exercise one's knowledge of an answer to the question by consciously thinking that answer, an episode of thinking that is partly constituted by the conscious, occurrent vehicle of that thinking—for example, an episode of one's consciously thinking that Fanny is bored.

The whole procedure would amount to intentionally answering one's question by intentionally bringing its answer to mind—in the example, by intentionally thinking that Fanny is bored. The intention expressed in that last act could not be a prior intention, on pain of one's needing to have brought the answer to mind in advance of intentionally bringing it to mind. For all that, the act of exercising one's knowledge or standing belief by thinking that Fanny is bored is performed intentionally and it is therefore an act of which one can have non-observational knowledge.¹⁹

Now all of this has taken place privately. However, knowing that one has answered a question for oneself by thinking its answer, and knowing what one has thereby thought, one may wish to share what has happened with others by publicising one's thinking. Here, one knows that one thinks (say) that Fanny is bored and one's task is to make available to others that one thinks that, typically by using

¹⁹ Even if we felt forced to deny that the act of thinking that Fanny is bored can be intentional, that need not exclude the possibility that one who performs that act thereby has non-observational knowledge of their doing so. For we might allow that non-observational knowledge of that act figures essentially in whatever it is that one does intentionally in this case—perhaps, for example, by bringing to mind the answer to one's question. For relevant discussion, see McDowell 2011.

appropriate linguistic means. What one intends to do here is to voice one's thought: knowing that one's thought is (say) that Fanny is bored, what one intends to do is to voice one's thought by voicing the thought that Fanny is bored. The function of voicing one's thought in this case is that of publicising a piece of prior thinking and thereby making that thinking recognizable to an audience. One's intention to voice this specific thought, and one's knowledge of what one will thereby do, or what one is thereby doing, are available to one in advance of completing the act of voicing. One's knowledge and intention is explanatorily prior to the act of voicing that one undertakes.

Suppose that the publicising form of voicing one's thoughts were the only form. In that case, application of Ryle's weaker symmetry claim might seem to support a purely recognitional model of our comprehending engagement with another's speech. Let me explain.

Where one's speaking is a means to publicising one's prior thinking, one's perspective on one's own speech is partly independent of one's perspective on one's own thinking. One is treating one's speech as, in the first instance, a means of making one's thinking recognisable to another. The aim of making one's thinking recognizable, and one's knowledge of what one will do in making one's thinking recognizable, is explanatorily prior to one's voicing one's thinking, and so one's ability to voice one's thoughts forms part of a larger structure embedding, in addition to the voicing, both one's prior thinking and the knowledge of what it will take for one to put it into words.

In seeking to apply Ryle's weaker symmetry claim, we are imagining that the same knowledge of what it is to voice one's thoughts figures in both knowing how to voice one's own thoughts and in being able to recognize another's voicing of their thoughts. In the case of knowing how to voice one's thoughts, we have seen that this knowledge embeds knowledge of what one thinks, and so knowledge of what one is, or will be, doing in voicing that thought. That knowledge, sustained by one's prior intention, intervenes between one's initial thinking and its outward expression, and is explanatorily prior to the latter action. Symmetry considerations therefore suggest an analogous structure in the case of the ability to recognize what thought another is voicing. Specifically, the audience's recognizing what thought another is voicing intervenes between their perceiving the speaker's voicing of a thought and the audience's thinking that thought for themselves. Recognition of what the speaker said would be the immediate upshot of perception

of their saying it. Now, we should allow that one's consciously recognizing that someone has voiced a thought—for example, the thought that Fanny is bored—requires one's consciously thinking that thought.²⁰ Still, one's thinking the thought is mediated by one's explanatorily prior recognition that the speaker has voiced the thought that Fanny is bored. That is to say, the presence of that instance of thinking is explained by the presence of the instance of knowing of which that thinking is an element. The symmetrical structure of origination and reception of the publicising form of voicing thoughts is represented in figure 1.

Speaker: Thinking that *p*

→ Knowing that they think that *p*, intending to voice that thought, and so knowing that they will be, or are, voicing the thought that *p*.

→ Voicing the thought that *p*

Audience:

→ Perceiving the voicing

→ Recognising, and so knowing, that the speaker voiced the thought that *p*

→ Thinking that *p*

Figure 1.

The basic function of the publicising form of voicing one's thoughts is to make recognizable to another what one thinks. The publicising form of voicing therefore seems to require sensitivity to others' observational perspectives on one's speech. Focusing on this form of voicing, then, might seem to support the idea that we saw Stanley needed, that voicing one's thoughts exploits knowledge about the perceptible features of speech that might sustain, or thwart, an audience's exercise of appropriate recognitional abilities. So, one potential diagnosis of Stanley's resistance to the idea of an ability simply to voice one's thoughts would be that it is based on the idea that the publicising form is the only form of voicing one's thoughts.

²⁰ See, e.g., Burge 2005: 174ff.

There is, however, a form of voicing one's thoughts which sits unhappily with the function of publicising. That form is thinking out loud, for example, in audible (or otherwise perceptible) soliloquy. The publicising model would suggest a view on which the function of the perceptible aspects of thinking out loud is to make recognizable one's thinking. Since thinking out loud can naturally take place without another as audience, its perceptible aspects need not be designed to make one's thinking recognizable to another. And so, a natural proposal would be that those perceptible aspects are designed to make one's thinking recognizable to oneself. However, the perceptible aspects of thinking out loud do not serve to make recognizable to oneself what one thinks. When one thinks out loud, one does not find out what one thinks by exercising a recognitional capacity on the deliverances of perception of the sounds or shapes that one thereby produces. Although one can find out what one thinks through thinking out loud, the way in which one does so is through knowing what one is doing when one thinks out loud. And one knows what one is doing when one thinks out loud through being responsible for what one is doing and the way in which that sustains one's having non-observational knowledge of what one is doing. On the face of it, then, thinking out loud seems to present a distinct form of voicing one's thoughts. Let's consider the activity of thinking out loud in more detail.

Thinking out loud is plausibly a basic form of voicing one's thoughts in which one's speaking (or signing, &c.) serves as a vehicle for one's thinking. We noted that occurrent thinking seems to require an occurrent vehicle. In the case of thinking in silent soliloquy, that vehicle is provided by conscious, occurrent aspects of one's thinking—for example, by one's imagining speaking. And I suggested that one way of exercising knowledge or dispositional belief would be through such occurrent thinking, for example in response to questions one has asked oneself. In originating such occurrent thinking, one thereby originates its occurrent vehicles. And originating such occurrent thinking is something that, in Hornsby's sense, one can simply do. The proposal now is that another way of exercising knowledge or dispositional belief through such occurrent thinking is by thinking out loud. In the case of thinking in perceptible soliloquy, the vehicle of one's thinking is the perceptible aspects of the act—for example, one's speaking. In originating occurrent thinking by thinking out loud, one thereby originates its occurrent vehicles, episodes of speaking. And originating such occurrent thinking in thinking out loud is something that one can simply do.

Thinking out loud differs from the publicising form of voicing one's thoughts in two main respects. First, in thinking out loud, by contrast with publicising, one's thinking depends constitutively on one's speaking. Indeed, both one's thinking and one's speaking plausibly depend constitutively on the more fundamental activity of thinking out loud.

One can be tempted into conceiving of thinking out loud as involving a sort of combination of two separable acts: an act of thinking; and an act of speaking. This is apt to seem plausible, as Matthew Soteriou has emphasised (2009, 2013: 238–246), because it is possible both to undertake the kind of speaking which figures in episodes of thinking out loud without thereby thinking and to undertake the kind of thinking which figures in episodes of thinking out loud without thereby speaking. However, as Soteriou also points out, it does not follow that acts of thinking and acts of speaking are more fundamental than acts of thinking out loud. Rather, the most fundamental act type here is thinking out loud, or thinking in, or by, speaking. Episodes of that kind must also be of the kind: thinking. And they must be of the kind: speaking. So, performing an act of thinking out loud will involve performing an act of thinking and an act of speaking. And acts of either kind could have been performed separately. However, the episode of thinking that figures in the episode of thinking out loud could not have occurred without the speaking; and that episode of speaking could not have occurred without the thinking. So, although episodes of the sub-kinds of act that figure in episodes of the kind thinking out loud can occur separately, the episodes of those kinds that figure in an episode of thinking out loud could not have occurred without so figuring. More is required of episodes of those kinds than mere co-occurrence, if they are to constitute an episode of thinking out loud. Those episodes must interlock through being elements in a more fundamental episode of thinking out loud. (Soteriou 2009, 2013: 238–246; Longworth 2008.)

Not only is the act of thinking out loud at least as fundamental as its constituent acts of thinking and speaking, but it is also something that one can do intentionally without separately intending to think or intending to speak. That is to say that, although one's thinking out loud depends on one's thinking and depends on one's speaking, one's thinking out loud is something that one can *simply* do.

The second main respect in which thinking out loud differs from the publicising form of voicing one's thoughts is that the speaking that figures in one's thinking out loud is most immediately

in the service of one's thinking, rather than in the service of making recognizable one's thinking. What one aims to do in thinking out loud is to engage in some thinking, typically thinking that meets some more or less generic specification. For example, one's aim might be to engage in thinking that answers the question, What is Fanny's state of mind? That specification of one's aim might or might not involve appeal to features of the sort of speaking that would be involved in so thinking. But the activity is aimed most fundamentally at thinking, or thinking out loud, in a way that satisfies the specification, rather than at thinking out loud in such a way as to make recognizable what it is that one thinks.

We sometimes characterise cases of this sort of thinking out loud as cases in which one doesn't know what one thinks until one hears what one says. And that characterisation might suggest that the perceptible aspects of thinking out loud do play an important role in making recognizable to one what it is that one thinks. However, being able to hear what one says in thinking out loud is most naturally thought of as, at most, an enabling condition for one's intentionally thinking out loud, where it is one's intentionally thinking out loud which explains one's knowing what one is doing, and so what one is thinking. The knowledge one has of what one is doing, and so what one is thinking, is plausibly non-observational knowledge. If that is right, then a better way of capturing what the characterization gestures towards would be this. One sometimes doesn't know what one thinks about some matter until one thinks it out loud, at which point one has distinctively practical, non-observational knowledge of what one doing, and so of what one thinks. But we can allow that our thinking out loud, and so our knowing what we are doing and what we are thinking, is sometimes enabled by observational knowledge of what we are saying.²¹

Thinking out loud is something that one can simply do which involves thinking with a perceptible vehicle. What does Ryle's weaker symmetry claim suggest about an audience's mode of engagement with someone's thinking out loud?

The most significant difference between thinking out loud and publicising is that in the former, unlike the latter, the speaker's voicing their thinking explains their knowing what they are thinking. As we saw, in the publicising form, the speaker can know in advance of voicing their thought what their doing so will involve; by contrast, in the thinking out loud form, the voicing is explanatorily prior to their knowing what they are doing. That is

²¹ For further discussion and defence of this aspect of thinking out loud, see Roessler 2015.

so, in turn, because the piece of occurrent thinking which enables them to know what they think, and so what voicing their thought will involve, and their knowledge of what they think have the same occurrent vehicle, the episode of speech in which they voice their thought. Considerations of symmetry therefore suggest that the audience's knowledge of what the speaker is doing in voicing their thought will also be an aspect of an episode of the audience's own thinking of that thought. That marks a point of similarity with publicising. However, in the case of publicising, symmetry considerations indicate that the audience's knowledge is like the speaker's in being explanatorily prior to their thinking. In the case of thinking out loud, by contrast, symmetry considerations suggest either that the audience's knowledge and their thinking are, like the speaker's, explanatorily on a par or that the audience's thinking has explanatory priority. And the latter view, on which thinking has priority over knowing, seems to be supported by the distinctively receptive use to which the audience puts their knowledge of how to think out loud.

My suggestion is that the perceptible vehicle of someone's thinking out loud can serve a suitably equipped audience as the vehicle of their own thinking. It is this shared vehicle which is the locus of the integration of the speaker's thinking with the audience's, integration which is required for their thinking to be collective. For example, consider a case in which someone simply thinks out loud that she is bored, by exploiting as perceptible vehicle an utterance of the sentence, "I am bored." Someone with the same ability simply to think out loud would thereby be placed in a position to engage in an occurrent form of thinking that the speaker is bored involving, as its vehicle, the *speaker's* utterance of the sentence, "I am bored." Like the speaker's knowledge of what they are thereby doing, the audience's knowledge of what they are doing in thinking that thought would arise with, or from, their thinking, rather than being explanatorily prior to it.

The core idea here is that knowing how to think out loud incorporates knowing how to think with a specific kind of perceptible vehicle. And knowing how to think with a specific kind of perceptible vehicle can give someone the cognitive wherewithal not only to think out loud, through the origination of such a perceptible vehicle, but also to think receptively, though the perception of a vehicle of thinking that was produced by someone else as a constituent of their own thinking. Crucially, considerations of symmetry in the speaker's and the audience's knowledge of how to think out loud suggests that, in both of their cases, that their

knowledge of what they are doing is explanatorily coeval, or posterior to, the thinking that they undertake through their respective engagements with the speaker's speaking. The speaker's speaking serves as the locus for integrating the speaker's and the audience's thinking and so as underpinning their collective thinking. At the same time, that shared vehicle figures in constituting their knowledge of what they are doing: the speaker's knowledge of the thought they are voicing; and the audience's knowledge of the thought that the speaker is voicing and that they have been brought to think receptively through their engagement with the speaker's voicing of it. The situation is represented in figure 2.

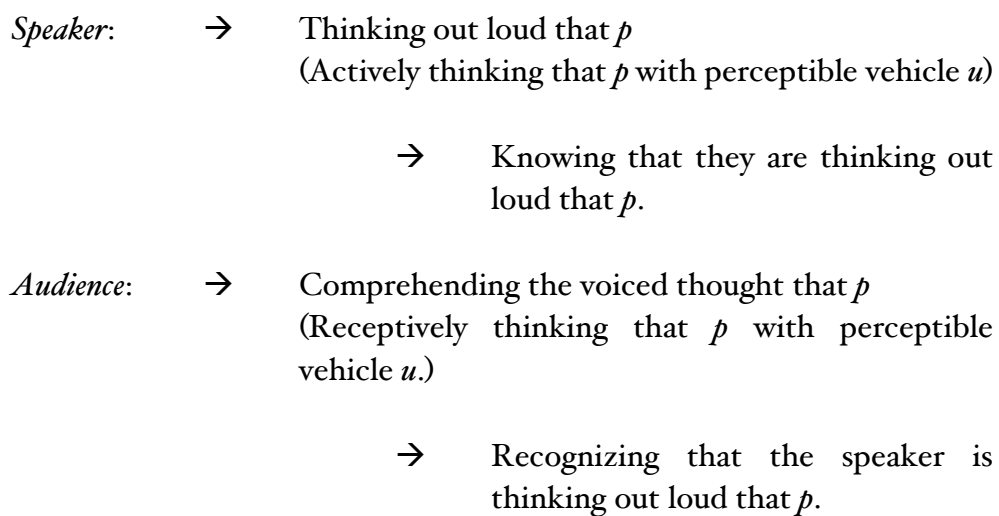


Figure 2.

In §3, we saw Dummett agree with Ryle's strong symmetry claim in holding that being able to comprehend others' speech depends on knowing how to speak. Let me conclude this section by noting a further convergence with Dummett, this time connecting the present application of Ryle's weaker symmetry claim with the following:

By the very nature of language, we could not learn its use as a means of interacting with others without simultaneously learning to use it as a vehicle for our own thoughts. (Dummett 1991: 103)

7. We began by considering Fanny Seward's view that her mother was someone with whom she thinks rather than (merely) speaks. I suggested that a full account of bare thinking together will need to find space for the possibility of thinking determinate thoughts

together. We have now considered two possible forms that might be taken by thinking determinate thoughts together: one involving mutual engagement with instances of the publicising form of voicing thoughts; and a second involving mutual engagement with instances of thinking out loud. I think that it is clear that the second form represents the purer example of thinking a determinate thought together, as opposed to collectively exercising knowledge of a speaker's occurrent thinking. For in the thinking out loud form, the speaker's and audience's thinking is integrated directly via sharing a vehicle, rather than being integrated only indirectly via shared knowledge of what the speaker is doing. However, my main purpose here has been to make space for any such possibility, with further articulation of that space being of secondary importance.

A further aspect of Seward's view that we considered was that thinking together is a distinctive achievement, demanding of special intimacy. Here, we have appealed to the idea that the operative demands on speakers and audiences concern the possession by those able to think together of the very same forms of know how. To think determinate thoughts together with someone, you and they must possess and exercise the same cognitive wherewithal to voice your thoughts. The requirements of sameness here can be met at increasingly demanding levels of fineness of grain and achieving convergence at the more demanding levels plausibly requires mutual harmonization over extended periods of interaction. And so, although we might allow, for certain purposes, that strangers or mere acquaintances can think together, the collective ability to simply think together, and to do so in ways that enforce matching amongst the precise thoughts engaged by that thinking, might well require more.²²

References

- Anscombe, G. E. M. 1957. *Intention*. Oxford: Basil Blackwell.
- Baier, Annette C. 1997. "Doing Things with Others: The Mental Commons." In L. Alanen, S. Heinamaa, and T. Wallgren eds. *Commonality and Particularity in Ethics*. London: Palgrave Macmillan: 15-44.
- Burge, Tyler. 2005. *Truth, Thought, Reason*. Oxford: Oxford University Press.

²² For discussions of an especially demanding case, that of first-personal thoughts, see Longworth 2013, 2014

- Butterfill, Stephen A. and Sinigaglia, Corrado. 2022. "Towards a Mechanistically Neutral Account of Acting Jointly: The Notion of a Collective Goal." *Mind*.
- Dummett, Michael. 1975. "Frege's Distinction Between Sense and Reference." In his *Truth and Other Enigmas*. London: Duckworth, 1978: 116-144.
- Dummett, Michael. 1991. *The Logical Basis of Metaphysics*. Cambridge, Mass.: Harvard University Press.
- Heck, Richard Kimberly. 1995. "The Sense of Communication." *Mind*, 104, 413: 79-106. (Originally published under the name "Richard G. Heck, Jr".)
- Heck, Richard Kimberly. 2006. "Reason and Language." In C. Macdonald and G. Macdonald eds. *McDowell and His Critics*. Oxford: Blackwell: 22-45. (Originally published under the name "Richard G. Heck, Jr".)
- Higginbotham, James. 1983. "The Logic of Perceptual Reports: An Extensional Alternative to Situation Semantics." *Journal of Philosophy* 80, 2: 100-127.
- Hornsby, Jennifer. 2005. "Semantic Knowledge and Practical Knowledge." *Proceedings of the Aristotelian Society, Supplementary Volumes*, 79: 107-130.
- Longworth, Guy. 2008. "Comprehending Speech." *Philosophical Perspectives* 22: 339-373.
- Longworth, Guy. 2013. "Sharing Thoughts About Oneself." *Proceedings of the Aristotelian Society*: 113: 57-81.
- Longworth, Guy. 2014. "You and Me." *Philosophical Explorations*, 17, 3: 289-303.
- Longworth, Guy. 2018. "Understanding What was Said." *Synthese* 195: 815-834.
- Longworth, Guy. 2019. "Sharing Non-Observational Knowledge." *Inquiry*: 1-22.
- McDowell, John. 1994. "Knowledge by Hearsay." In B. K. Matilal and A. Chakrabarti eds. *Knowing from Words*. Dordrecht: Springer: 195-224. Reprinted in his 1998: 414-443.
- McDowell, John. 1998. *Meaning, Knowledge, and Reality*. Cambridge, Mass.: Harvard University Press.
- McDowell, John. 2011. "Anscombe on Bodily Self-Knowledge." In A. Ford, J. Hornsby, and F. Stoutland eds. *Essays on Anscombe's Intention*. Cambridge, Mass.: Harvard University Press: 128-146.
- Oliver, Alex and Smiley, Timothy. 2013. *Plural Logic*. Oxford: Oxford University Press.

- Roessler, Johannes. 2015. "Self-Knowledge and Communication." *Philosophical Explorations* 18, 2: 153–168.
- Small, Will. 2019. "Basic Action and Practical Knowledge." *Philosophers' Imprint* 19, 19: 1–22.
- Soteriou, Matthew. 2009. "Mental Agency, Conscious Thinking, and Phenomenal Character." In L. O'Brien and M. Soteriou eds. *Mental Actions*. Oxford: Oxford University Press: 231–252.
- Soteriou, Matthew. 2013. *The Mind's Construction: The Ontology of Mind and Mental Action*. Oxford: Oxford University Press.
- Stanley, Jason. 2005. "Hornsby on the Phenomenology of Speech." *Proceedings of the Aristotelian Society, Supplementary Volumes*, 79: 131–145.