Episteme (2026), **23**, 173–194
doi:10.1017/epi.2026.10101

episteme

ARTICLE

Achieving Common Knowledge

Thomas Lockhart

Philosophy, Auburn University, USA
Email: lockhartt@auburn.edu

(Received 8 June 2025; revised 7 November 2025; accepted 27 November 2025; first published online 02 March 2026)

Abstract

Accounts of common knowledge typically focus on how to represent what agents would know when they achieve common knowledge, instead of asking how agents achieve common knowledge. In this paper, by modifying a definition proposed by David Lewis, I generate a set of conditions the satisfaction of which would entail that a group of agents has common knowledge, properly so-called. I then ask what it would take to satisfy these conditions. I argue that these conditions can be satisfied and common knowledge achieved in the course of a conversation if, but only if, we adopt a particularly strong interpretation of Austin's claim that the successful performance of a speech act requires the securing of uptake. I then argue that the generalization of this result to conversations involving more than two people requires that the audience to whom a speech act is addressed be able to recognize themselves as the audience of the speech act, in an essentially plural act of recognition. I close by arguing that one of the so-called paradoxes of common knowledge (the coordinated attack problem) in fact gives us some reason to favor Austin's account of speech acts and therefore some reason to believe that common knowledge can, on occasion, arise.

Keywords: Common knowledge; JL Austin; illocution; David Lewis; coordinated attack

1. Introduction

Can it ever come about that some proposition C is common knowledge for some group of individuals, G? For the sake of this paper, I assume that when C is common knowledge for a group of individuals G, it is true that: everybody in G knows that C, everybody in G knows that everybody in G knows that C, and so on. I call this infinite hierarchy the CK hierarchy. So, put more specifically, the question I address in this paper is: Can the CK hierarchy ever come about?

In order to address this question, I propose (Section 3) a tractable set of conditions the satisfaction of which would entail the CK hierarchy. I develop this set of conditions by modifying David Lewis' original account of common knowledge (which I describe in Section 2). Lewis' conditions must be modified for my purposes because he does not, in fact, define what we might think of as common knowledge, properly so-called. Instead of "knows," Lewis operates with a

© The Author(s), 2026. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0/), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

weaker notion, that of "reason to believe." Having modified Lewis' account of common knowledge, I then ask what it would take to satisfy this set of conditions (Section 4). It turns out that what are typically taken to be paradigmatic examples of common knowledge – cases in which some everyday object is publicly perceptible by a group of individuals, each of whom has ordinarily functioning sensory equipment – do not appear to satisfy these modified conditions. But if not this kind of paradigmatic case, what could constitute an example of common knowledge, properly so-called?

My claim in Section 4 is that if we adopt a particular interpretation of Austin's claim that the performance of a speech act involves the securing of what he calls "uptake," then it will be possible to achieve common knowledge properly so-called in a conversation between two people. I further argue that if Austin is wrong, then it is impossible to achieve such conversational common knowledge. I then consider (Section 5) what would be required to extend this result to a conversation involving more than two people. I argue that the conditions for genuine common knowledge in such a case are in fact more demanding: for common knowledge to come about in the course of a conversation involving more than two people, we need to insist that in order for the performer of the illocution to secure uptake of that illocution, the audience which is the target of the speech act must recognize *themselves* as the audience of the speech act, where the *themselves* cannot be understood severally, but must be understood collectively; that is *themselves*, as used here, is irreducibly plural.

This argument does not, however, establish that such common knowledge in fact ever comes about. I close (Section 6) by suggesting that one of the so-called paradoxes of common knowledge – the coordinated attack problem – counts in favor of thinking that common knowledge can and does occur in the course of conversation.

2. Lewis' account

My question is whether the CK hierarchy could ever come about. There are various accounts of common knowledge available. These accounts differ in precisely how they define common knowledge, although typically it can be shown that the CK hierarchy follows from each definition of common knowledge. However, few of these definitions are couched in such a way that it will be possible to ask how it could be that a group of agents could come to possess common knowledge.[1] The main exception is the definition proposed by Lewis. As Cubitt and Sugden put the point, Lewis' account "doesn't simply *represent* what individuals" would know, once they had achieved common knowledge, but instead seeks to "*explain* how those individuals come to know what they are represented as knowing" (Cubitt and Sugden 2003: 207).[2] However, Lewis' definition does not, in fact, generate the CK hierarchy itself, but a slightly different infinite hierarchy. In this Section, I'll explain Lewis' account. In Section 3,

[1]Examples include Aumann (1976), Barwise (1988), Binmore and Branderburger (1988), Monderer and Samet (1989). See also Vanderschraaf and Sillari (2023) for discussion.

[2]Another account, which is exceptional in this regard, is that proposed by Gilbert (1989: 188–189), which certainly appears to be a set of conditions explaining how a group of agents could come to possess common knowledge. Briefly, my reasons for not using Gilbert's account are these. First, one of the clauses of her account is: "A, B, and C perceive each other perceiving that p" (Gilbert 1989: 189). But what is it to perceive another perceiving that p? In particular, if I perceive another perceiving that p, does that put me in a position to know that they know that p? Not obviously, for couldn't somebody perceive that p but fail to come to know that p? Second, Gilbert's account doesn't, in fact, generate the CK hierarchy. The problem is that perceiving each other, perceiving that p can generate only very few levels of the hierarchy. For it can be the case that A perceives B perceiving that p, and B perceives A perceiving that p, without it being the case that B knows (or perceives) that A perceives B perceiving that p, and without it being the case that A knows (or perceives) that B perceives A perceiving that p.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

I explain why it needs to be modified. In the following Sections, I'll turn to the question of what it would take to satisfy this set of modified conditions.

I emphasize at the outset that it might be possible for common knowledge properly so-called to come about even though the modified set of conditions I extract from Lewis is not satisfied. My claim is not that the definition I give of Strong Common Knowledge is *the* correct such definition, both necessary and sufficient for common knowledge, properly so-called. The question animating my inquiry is simply whether we can find any situations in which the CK hierarchy holds. I observe that Lewis' account, when modified, yields a tractable set of conditions the satisfaction of which will entail the CK hierarchy. And I then claim that there are certain cases involving speech acts which (under certain assumptions about the nature of speech acts) satisfy the modified Lewisian conditions I propose, and which therefore would (if those assumptions were correct) count as cases of common knowledge, properly so-called.

I turn now to Lewis' account. According to Lewis, C is (what he calls) "common knowledge" when:

Lewisian common knowledge

LCK-1 Everyone in G has reason to believe that A holds.

LCK-2 A indicates to everyone in G that everyone in G has reason to believe that A holds.

LCK-3 A indicates to everyone in G that C.[3]

"We can", Lewis says, "call any such state of affairs A a *basis*" for common knowledge in G that C (Lewis 1969: 56).

In a moment, I'll explain how Lewis' definition of common knowledge works. I start with three prefatory observations. First, Lewis' definition uses "reason to believe." "Reason to believe" is, of course, not the same as "knowledge." It is ultimately for this reason that Lewis' account will need to be modified in order to characterize common knowledge, properly so-called.

Second, Lewis' definition relies on the notion of *indicating*:

Indicating A *indicates* to someone X that Δ if and only if, if X had reason to believe that A held, X would thereby have reason to believe that Δ.[4]

The definition of indicating is indexed to an individual: it may well be that what some particular A indicates to Gwendolen is different from what that particular A indicates to Cecily if Gwendolen and Cecily happen to have different "inductive standards and background information" (Lewis 1969: 53). Crucially, in any particular case, Lewis assumes that the members of G have "the same inductive standards and background information, at least nearly enough so that A will indicate the same things" to all the members of G (Lewis 1969: 53).

The definition of indicating also makes use of the word "thereby." This is, as will become clear below, essential.[5]

[3]This definition is at Lewis (1969: 56), I've modified only the variable names.

[4]This is at Lewis (1969: 52–53), I've modified only the variable names.

[5]Cubitt and Sugden also draw attention to the importance of the word "thereby" for Lewis. They suggest that when A indicates Δ to X, X's "reason to believe that A holds *provides [X's] reason for* believing that [Δ] is true" (Cubitt and Sugden 2003: 186) The indication relation will thus employ not just the principles of deductive reasoning, but also those other inference rules proper to the "logic of reasoning [X] endorses" (Cubitt and Sugden 2003: 187).

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

Third, the notion of a Lewisian basis is tractable. It is not, as we shall see, straightforward to understand. Nevertheless, it is possible to check, for any state of affairs, whether it counts as a Lewisian basis for common knowledge, as he defines that term.

In order to see how Lewis' definition of common knowledge works, we can start with an example. Suppose that Gwendolen and Cecily are out hiking together. A large bear slowly ambles onto the trail ahead, follows the trail for a little while, then heads back into the woods. Wordlessly, Gwendolen and Cecily stop to watch the progress of the bear. Set:

(A–1) Cecily and Gwendolen are hiking together. They both stop when a large bear appears and slowly follows the trail for a while before heading into the woods.
(C–1) There is a bear nearby.

Now, let us see why A–1 is a basis for C–1 to be what Lewis calls "common knowledge."

To see that LCK-1 holds, we must assume that each of the members of G has an ordinarily functioning perceptual apparatus, so let us make this assumption. We'll also need to assume that Gwendolen and Cecily have relatively ordinary inductive standards and background information, so that bears are familiar enough animals that both Gwendolen and Cecily can readily identify them.

What about LCK-2? Since there are two individuals in our scenario, we need to check four things:

2a. If Gwendolen has reason to believe that A–1 holds, then Gwendolen thereby has reason to believe that Cecily has reason to believe that A–1 holds.
2b. If Gwendolen has reason to believe that A–1 holds, then Gwendolen thereby has reason to believe that Gwendolen has reason to believe that A–1 holds.
2c. If Cecily has reason to believe that A–1 holds, then Cecily thereby has reason to believe that Gwendolen has reason to believe that A–1 holds.
2d. If Cecily has reason to believe that A–1 holds, then Cecily thereby has reason to believe that Cecily has reason to believe that A–1 holds.

Now, because we are assuming that Gwendolen and Cecily are likewise situated with respect to the relation of indication, whatever reasoning can convince us that 2a is true will, mutatis mutandis, convince us that 2d is true, and vice versa (and likewise for 2b and 2c). Furthermore, the description I have given of A–1 shows that Gwendolen and Cecily each stand in the same epistemic relation to A–1 as does the other. So, given that they have the same inductive standards and background information, whatever reasoning can convince us that 2a is true will also convince us that 2c is true, and vice versa (and likewise for 2b and 2d).[6] So, in this particular case, we can convince

[6]One might worry that if something like the KK principle is true, then it will be trivially easy to satisfy 2d: any case in which Cecily has reason to believe A–1 would seem, via an application of the KK principle, to be a case in which Cecily has reason to believe that Cecily has reason to believe that A–1. And if 2d can be satisfied this easily, then it ought not entail 2a, as I have claimed, for 2a is not trivially true. It is at this point that the word "thereby" is crucial: when 2d is true, it must be in virtue of the content of A–1 (if we follow the interpretation of "thereby" suggested by Cubitt and Sugden (see footnote 5, above) Cecily's reason for believing A–1 must provide her reason for believing that she has reason to believe A–1). In other words, the KK principle can make

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

ourselves that all of 2a–2d are true simply by convincing ourselves that one of them is true.[7]

To see that 2a holds, we need to show that if Gwendolen has reason to believe that A–1 holds, then she thereby has reason to believe that Cecily has reason to believe that A–1 holds. It is plausible to think that this is so. For if Gwendolen has reason to believe that Cecily was hiking with her and they both stopped upon the appearance of the bear, then Gwendolen will thereby have reason to believe that Cecily will also (a) have reason to believe that they were both hiking together but stopped, and (b) have reason to believe that a bear just up ahead followed the trail for a while.

LCK-3 is, on the assumption that Cecily and Gwendolen both have relatively ordinary inductive standards and background knowledge, also straightforward. What we need to see is (a) if Cecily has reason to believe that A–1 holds, she would thereby have reason to believe that there is a bear nearby and (b) if Gwendolen has reason to believe that A–1 holds, she would thereby have reason to believe that there is a bear nearby. Now, since we are assuming that the indication relation is the same for each of them, then we only need to check (a). And (a) is surely a reasonable inference for most people. I, therefore, propose to stipulate that Cecily's background knowledge and inductive standards are such as to ensure the satisfaction of LCK-3.

Lewis goes on to show that the satisfaction of the clauses of Lewisian Common Knowledge will generate what we can call the RTB hierarchy: both Gwendolen and Cecily have reason to believe C–1, both Gwendolen and Cecily have reason to believe that both Gwendolen and Cecily have reason to believe C–1, and so on. I omit the details, since these can be found elsewhere (Cubitt and Sugden 2003: 204–206). But it is worth identifying the central ideas of Lewis' demonstration. Lewis establishes a key lemma:

Lemma. *If (i) A indicates to X that Y has reason to believe that A, and if (ii) A indicates to X that* Δ*, and if (iii) X has reason to believe that Y shares X's inductive standards and background information, then A indicates to X that Y has reason to believe that* Δ.

Recall that LCK-2 states that "A indicates to everyone in G that everyone in G has reason to believe that A holds." This is an instance of clause (i) of the lemma. And LCK-3 states that "A indicates to everyone in G that C." This is an instance of clause (ii). So, we can use the lemma to combine LCK-2 and LCK-3 to yield:

IND-1 A indicates to everyone in G that everyone in G has reason to believe that C.

But now, IND-1 is itself an instance of clause (ii) of the lemma. So, we can use the lemma to combine LCK-2 with IND-1 to yield:

(2d′) If Cecily has reason to believe that A–1 holds, then Cecily has reason to believe that Cecily has reason to believe that A–1 holds.

true, but not 2d.

[7]It will not, in general, always be true that we need check only one case. The various members of group G might stand in different epistemic positions with respect to the situation described in the putative basis for common knowledge.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

IND-2 A indicates to everyone in G that everyone in G has reason to believe that everyone in G has reason to believe that C.

We can continue in this way indefinitely, at each step applying LCK-2 and Lewis' lemma to the previous result.

At the heart of this indefinite series of inferences is the fact that when A is a basis, it has the feature that A indicates to everyone in G that everyone in G has reason to believe that A itself holds. In other words, it is true of every individual X in G that if X had reason to believe that A holds, then X would thereby have reason to believe that everyone in G has reason to believe that A holds. To put the point yet another way: A is such that one of the things one has reason to believe, when one has reason to believe A, is that all the members of G also have reason to believe A.

3. A modification to Lewis' account

Lewis' account of common knowledge works with the notion of "reason to believe." "Reason to believe" is not the same as "knowledge," and it is not straightforward to see how to convert a reason to believe *p* into knowledge that *p*. In fact, as I understand Lewis' account, C can be (what he calls) common knowledge without any of the members of G knowing C, or even believing C. How might this come about?

One might think the answer is straightforward: somebody can have reason to believe something, but, being finite and busy, simply not take the time to come to believe it. But this is not my worry about Lewis' use of "reason to believe." I think that it is possible to believe infinitely many things without having taken time to think about each of them. I'll return to this point briefly at the end of this Section. Before that, I describe why the notion of "reason to believe" doesn't give an adequate analysis of the notion of common knowledge.

To see the problem, consider the following case: Gatsby has offered to have Nick's grass mown because Gatsby is planning to meet Daisy at Nick's house. Gatsby arrives for the meeting, and Nick narrates the following exchange between himself and Gatsby, who speaks first:

> "Is everything all right?" he asked immediately.
> "The grass looks fine, if that's what you mean."
> "What grass?" he inquired blankly. "Oh, the grass in the yard." He looked out the window at it, but, judging from his expression, I don't believe he saw a thing.
> "Looks very good," he remarked vaguely. (Fitzgerald 2025: ch V)

As I understand the case, Nick knows perfectly well that Gatsby can distinguish between a well-mown lawn and an overgrown lawn. But although Gatsby is looking out the window at the lawn, and although Gatsby even *says* that the lawn "looks very good," "That the lawn looks good" is not common knowledge between Nick and Gatsby, for Nick has reason to doubt that Gatsby has, in fact, seen the lawn and determined that it looks good.

First, consider: does Gatsby have reason to believe that the lawn looks good? We know that he cares about the appearance of the lawn, we assume that he is a competent judge of such things, and he is looking directly at the lawn. Crucially, the situation seems precisely parallel to the situation described in A–1 (Cecily, we've assumed, has the ability to identify bears, and she stops and looks directly at one on the path ahead). Just as we

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

said that Cecily has reason to believe A–1, my own inclination would be to say that Gatsby does indeed have reason to believe that the lawn looks good.[8]

But now consider Nick's situation. Does Nick have reason to believe that Gatsby has reason to believe that the lawn looks good? I've just suggested that because Cecily and Gatsby are in parallel situations, *we* have reason to believe that Gatsby has reason to believe that the lawn looks good. But all of this information about Gatsby is also available to Nick. So, my own inclination is to think that Nick does indeed have reason to believe that Gatsby has reason to believe that the lawn looks good. However, it would be a mistake for Nick to come to believe that Gatsby believes that the lawn looks good, for Nick also has reason to believe – in this case, stronger reason to believe – that Gatsby hasn't paid any attention to the lawn.

I just claimed that Nick has reason to believe that Gatsby has reason to believe that the lawn looks good. But if this is right, then ultimately the Gatsby case is parallel to A–1, and Lewis' account of common knowledge would predict that "The grass looks good" is common knowledge between Nick and Gatsby.[9] But this can't be right, because Nick knows perfectly well that Gatsby doesn't believe that the lawn looks good. The problem is that on the way I've proposed to interpret "reason to believe," one can have reason to believe p and yet also recognize that one has better reason not to believe p.

One could certainly try to dodge this problem by arguing that my preferred way of understanding "reason to believe" is incorrect. One could insist that we must understand "reason to believe," so that it will be true to say that Gatsby does not have reason to believe that the grass looks good, and Nick does not have reason to believe that Gatsby has reason to believe that the grass looks good. Call this the stronger interpretation of "reason to believe."

The problem now is that if we adopt the stronger interpretation of "reason to believe", my argument (from Section 2) that A–1 can serve as a basis for Lewisian common knowledge that C–1 is rendered implausible. Just as Gatsby can look at the lawn without seeing it, surely we can imagine Cecily, in A–1, being sufficiently distracted such that she can look at the bear without seeing it. A–1, as described, tells us nothing to rule out the possibility that she is thus distracted. If we insist on ruling out the possibility of such a distraction in order for it to be true to say that Cecily has reason to believe that there is a bear on the path, then A–1 is not, as it stands, a Lewisian basis for common knowledge.

Neither of these choices seems satisfactory. A–1 seems plausible as a Lewisian basis for common knowledge if we interpret "reason to believe" in my preferred sense. But then, we cannot rule out the Gatsby case, which is clearly not a case of common knowledge. But if we interpret "reason to believe" more strongly, A–1 seems significantly less plausible, as it stands, as a basis for Lewisian common knowledge.

I suggest, therefore, that "reason to believe" does not yield an account of common knowledge. I propose instead that we shift attention directly to the notion of knowledge. Can we somehow strengthen Lewis' account so as to directly yield the CK hierarchy?

The modification I propose is relatively simple: to replace "reason to believe" with "knows" in Lewis's account of common knowledge.

[8]One option for Lewis here would be simply to disagree and to say that Gatsby doesn't have reason to believe that the lawn looks good. Another possibility would be to invoke the following principle: "Anyone who has reason to believe something will come to believe it provided he has a sufficient degree of rationality" (Lewis 1969: 55). On this approach, Lewis could say that the reason Gatsby hasn't come around to believing that the lawn looks good is a want of rationality, for his attention is focused elsewhere.

[9]More precisely, "Nick and Gatsby are together looking out the window at the freshly mown grass" would be the basis for common knowledge between them that "The lawn looks good".

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

Strong common knowledge

SCK-1 Everyone in G knows that A holds.
SCK-2 A K-indicates to everyone in G that everyone in G knows that A holds.
SCK-3 A K-indicates to everyone in G that C.
K-Indicating A *K-indicates* to someone X that Δ if and only if, if X knew that A held, X would thereby know that Δ.

We could then derive the CK hierarchy in essentially the same way that Lewis derives the RTB hierarchy, changing "reason to believe" to "knows" as appropriate.[10]

One might object that it is simply psychologically impossible for the CK hierarchy ever to come about, for the CK hierarchy attributes to each subject in G infinitely many pieces of knowledge. I concede that if you think that, in order for *S* to know *p*, *S* must have taken time to contemplate *p*, or to have taken time to infer *p* from her stock of existing knowledge, then you must reject the CK hierarchy. However, as Lederman points out, there is a familiar use of the word "know" (and indeed "believe") according to which I know and believe infinitely many things:

> For all $n > 0$, I know that there are not exactly n unicorns. I know an infinite set of propositions concerning the nonexistence of unicorns; any theory of belief and knowledge must account for this datum. The way in which subjects who had common knowledge and common belief would know or believe infinitely many propositions does not appear to be importantly different from the way in which I know infinitely many propositions about the nonexistence of unicorns. (Lederman 2018b: 217).

Of course, there are many questions about how an account of belief and knowledge adequate to account for the fact that I know infinitely many propositions about unicorns should go. But I believe that nothing I say in what follows relies on any philosophical claims stronger than those that would be required for such an account.

4. Strong common knowledge

My question now is: Can we find some state of affairs which will constitute a basis for Strong Common Knowledge?[11] Stock examples of common knowledge tend to involve cases in which something is publicly perceptible to a group of individuals, all of whom have the standard perceptual apparatus. For example:

(A–2) Gwendolen and Cecily are hiking together. A bear appears on the trail ahead.
(C–2) There is a bear up ahead.

[10]See, for example, Vanderschraaf and Sillari (2023: §2.2). Such a derivation would also require us to assume that the members of G have common knowledge of each others' rationality. I propose to make this assumption. I note that Lederman (2018a) argues that it is an implausible idealization; I also note that Lewis writes "A basis for common knowledge generates higher-order expectations with the aid of pre-existing higher-order expectations of rationality. Can these themselves be generated by some basis for common knowledge? Yes, because all the higher-order expectations of rationality needed to generate an *n*th-order expectation are themselves of less than *n*th-order" (Lewis 1969: 57). The key 'inductive' step in the argument will require us to establish a version of Lewis' lemma, modified so as to concern knowledge.

[11]Thanks to an anonymous referee for helping me with the argument of this section.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

Although this case might seem to be paradigmatic, it does not satisfy our characterization of Strong Common Knowledge. We've already seen (in Section 3) the essential reason why this is so. To establish SCK-2, we'll need to show, among other things, that:

(1) If Gwendolen knows that A–2 holds, then she thereby knows that Cecily knows that A–2 holds.

But we cannot establish (1). First, observe that Gwendolen might not know that Cecily knows what a bear is, or how to identify them (this problem becomes more pressing if we replace "bear" with, for example, "kingsnake"). But even if we can fix this (perhaps by appropriately modifying how we handle our background assumptions), there is a second, more pressing, problem. It is perfectly possible that, on this particular occasion, Cecily has simply failed properly to exercise her perceptual capacity to recognize bears: even though it is in plain sight, on this occasion, she just hasn't recognized it. So, although she has reason to believe there is a bear on the path, she hasn't seen it, and so doesn't know it is there.

I am not here claiming that publicly perceptible facts can never yield an example of Strong Common Knowledge, as I have defined it. Perhaps there is some way that this can come about. My claim is simply that it is not at all straightforward to see how this could be. Nor am I claiming that publicly perceptible facts can never yield the CK hierarchy. For perhaps there is some other way to specify a tractable set of conditions the satisfaction of which will yield the CK hierarchy. My claim is simply that publicly perceptible facts do not furnish straightforward examples of Strong Common Knowledge.

Lewis himself suggests a more promising place to look for clear examples of Strong Common Knowledge. He asks us to consider a situation in which two people have been talking. Ernest tells Cecily that he intends to return to the same place tomorrow to continue the conversation, and as a result, Cecily comes to expect that Ernest will return.[12] So, let us set:

(A–3) Ernest and Cecily have been talking. Ernest tells Cecily that he intends to return to the same place at the same time tomorrow to continue the conversation.
(C–3) Ernest intends to return to the same place at the same time tomorrow to continue the conversation.

I'll argue that, on certain assumptions, A–3 can satisfy clauses SCK-1 and SCK-2 of the definition of Strong Common Knowledge. The A–3/C–3 pair does not satisfy SCK-3 as they stand, but I'll fix this at the end of the section by modifying the example.[13]

The assumptions we need come from Austin. Recall that Austin argues that the performance of an illocutionary act (such as telling) involves the securing of what he calls "uptake":

[12]This example is Lewis' (1969: 52), I've added names.

[13]Lederman (2018b) constructs an ingenious argument designed to show that common knowledge cannot come about in any circumstances. But Lederman's argument deals only with the kinds of perceptual cases I've just set aside. All of his examples involve situations in which the putative item of common knowledge is publicly perceptible, and his argument exploits the fact that anything publicly perceptible can appear slightly differently to one observer than to another. But the examples of Strong Common Knowledge I propose in what follows are not publicly perceptible. My examples involve conversational participants who have common knowledge about some proposition available to them on the basis of their participation in that conversation. I don't believe Lederman's argument applies to the cases I describe.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

> Unless a certain effect is achieved, the illocutionary act will not have been happily, successfully performed. This is not to say that the illocutionary act is the achievement of a certain effect. I cannot be said to have warned an audience unless it hears what I say in a certain sense. An effect must be achieved on the audience if the illocutionary act is to be carried out. How should we best put it here? And how can we limit it? Generally, the effect amounts to bringing about the understanding of the meaning and of the force of the locution. So, the performance of an illocutionary act involves the securing of *uptake*. (Austin 1962: 116–117)

I claim that if Austin, interpreted in a sufficiently strong way, is correct, then A–3 would satisfy clauses SCK-1 and SCK-2 of the definition of Strong Common Knowledge.

To show that A–3 satisfies SCK-1, we need to show that, on the assumption that A–3 is true, both Ernest and Cecily know that A–3 holds. So, what we need to show is:

i If A–3 is true, then Cecily knows it is true.
ii If A–3 is true, then Ernest knows it is true.

I assume that if Ernest and Cecily are engaged in a conversation, then they both know that this is so. The difficulty comes when we ask whether each knows that Ernest has told Cecily that he intends to return tomorrow to continue the conversation.

With respect to (i), observe first that if Austin is wrong about the nature of illocutionary acts, then (i) will not obviously hold. For if Austin is wrong, then Ernest can "tell" Cecily that *r* without uptake being secured. That is, Ernest can "tell" Cecily that *r* without Cecily knowing that Ernest has "told" her that *r*. But this means that A–3 could be true without Cecily knowing it to be true.

However, if Austin is right, then Ernest cannot have told Cecily *r* without Cecily "understanding the meaning and the force of the illocution." It is reasonable to think (and this is how I propose to interpret Austin's remark) that this means that Cecily will know that Ernest has told her (force) that he intends to return to the same place at the same time tomorrow to continue the conversation (meaning). Notice that in specifying what Cecily knows, I specified that *she'll* know that Ernest has told *her* such-and-such. It won't be enough for Cecily simply to know that some speech act of telling has occurred, and what the content of that speech act is. She'll need to know that it was she who was told, and by whom she was told. So, if Austin is right, (i) holds.

With respect to (ii), notice that if Austin is wrong, then (ii) is relatively easy to establish. For if Austin is wrong, then Ernest can "tell" Cecily that *r* even if uptake is not secured. He'll count as "telling" Cecily that *r* so long as his utterance has an appropriate meaning and he has the intention to bring about in Cecily an understanding of the force and content of his utterance. And it is reasonable to think that if he has done this, then he'll know that he has done so, as (ii) requires.

To establish that (ii) can be true if Austin is right to insist that performing an illocution requires the securing of uptake, we'll need it to be the case that when one secures uptake, one knows that one has done so. So, I'll make use of the following principle:

(2) If X intentionally told Y that *r*, then X knows that Y knows that Y has been told *r* by X.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

I think it is reasonable to think that (2) follows from Austin's account. For Austin, the successful performance of an illocutionary act involves "the securing of uptake". So, assuming that Ernest has intentionally[14] told Cecily that *r*, this will have required him intentionally to have secured in Cecily an understanding of the force and the content of his illocution. And this – intentionally securing in one's audience an understanding of the force and the content of one's illocutionary act – is not something one could do without knowing one has done so.

There is a natural way of attempting to resist the claim that when one intentionally performs an Austinian speech act, one does so knowingly. Since illocutions require uptake, one's attempts to perform an illocution will be unsuccessful if one fails to secure uptake. So, wouldn't it be more correct for Austin to say that all Ernest was knowingly able to do was to *try* to tell Cecily that *r*? But this objection can't be right as an interpretation of Austin, for it elides the distinction between illocutions and perlocutions. To see why, observe that the objection characterizes Ernest's action as an attempt to bring about a particular consequence. Whether or not his action succeeds depends on whether or not those consequences at which Ernest was aiming (viz., uptake) happen to come about. But this is a description of a perlocutionary act. The contrast between perlocutionary and illocutionary acts depends crucially on the fact that uptake is not a "downstream" consequence of the illocutionary act, but is instead a constitutive component of the act itself, and so must be part of what Ernest intentionally did.

The point here is contentious, both philosophically and as an interpretation of Austin. But since my goal is to try to describe a situation in which Strong Common Knowledge is possible, I'll simply record here the assumptions on which, I've claimed, we can find a case of Strong Common Knowledge.

Austin*:

a. Performing an illocutionary act requires securing in one's audience an understanding of the force and content of the illocution;
b. If X has intentionally performed some illocutionary act, then X knows that her audience knows that X has performed that illocutionary act with them as the audience.[15]

[14]The qualification is necessary because Austin acknowledges that one may on occasion unintentionally perform an illocutionary act (Austin 1962: 106, footnote). So it is possible that one might not know that one has performed an illocutionary act if the performance was unintentional. Since illocutionary acts require uptake, if X unintentionally performs some speech act then Y will understand the force and content of the illocution, but X won't know that Y knows this. So, assuming that one knows what one intentionally does, I require that X's performance was intentional in order to ensure that X knows that Y knows X has performed the speech act. It will follow from my adoption of (2) that my account will not apply to those illocutions which might possibly have been performed unintentionally; my account will apply only to those speech acts which could not possibly have been performed unintentionally, as I assume is the case with Ernest's speech act in A–3.

[15]One might claim that illocutionary acts, so described, are impossible. My claim is that if they are possible, they will provide an example of Strong Common Knowledge. I do not argue that they are in fact possible, although I take the considerations in Section 6 to give some evidence in favor of this possibility. Thanks to a referee for urging clarity on this point.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

I conclude that if Austin* is correct, (ii) will hold. So, I've established that if Austin* is correct, A–3 satisfies SCK-1.

I turn now to the claim that if Austin* holds, then A–3 will satisfy SCK-2. To show this, we'll need to establish the following:

i If Ernest knows that A–3 is true, then he thereby knows that Cecily knows that A–3 is true.
ii If Cecily knows that A–3 is true, then she thereby knows that Ernest knows that A–3 is true.[16]

To see that (i) is true if Austin* holds: in my discussion of SCK-1, I argued that in order for Ernest to tell Cecily that *r*, Ernest must see to it that she understands the force and the content of his illocutionary act, that is, he must see to it that she knows she has been told *r*. That means that one of the things he knows, on the assumption that he knows that he has told her that *r*, is that she knows that she has been told *r*. And this is what (i) requires.

To see that (ii) is true if Austin* holds: let us suppose that Cecily knows that A–3 is true. So, Cecily knows that Ernest has successfully performed a particular speech act, which, we are assuming, can only have been performed intentionally (see footnote 14). But if intentionally performed speech acts are performed knowingly, then for the speech act to have occurred, Ernest must know that he has successfully performed it. But this means that Cecily cannot know that the speech act has occurred if she doesn't know that Ernest knows he has successfully performed it, for to the extent that she harbors any doubt about whether Ernest knows he has succeeded, she must accordingly doubt whether the speech act has occurred. In other words, if Cecily does know that the speech act has occurred, then one of the things that she will thereby know is that Ernest knows he has performed it. And this is precisely what (ii) requires.

I've defended a conditional claim: if Austin* holds, then A–3 will satisfy SCK-1 and SCK-2. I emphasize that my claim is conditional. I have not taken it upon myself to give a full defense of Austin*.

SCK-3 is more troublesome. Even if Austin* is correct, Ernest can succeed in telling Cecily that he intends to return tomorrow without Cecily coming to believe him. I'll improve the example in a moment.

Austin's claim (let alone Austin*) is controversial.[17] If Austin's opponents are correct, then one can successfully perform an illocutionary act without securing uptake. And if this is so, then I see little hope that A–3 could be a basis for Strong Common Knowledge.

[16]Officially, there are two further claims we'd also need to establish:

iii If Ernest knows that A–3 is true, then he thereby knows that he knows that A–3 is true.
iv If Cecily knows that A–3 is true, then she thereby knows that she knows that A–3 is true.

But (as we saw in Section 2), because we are assuming that the indication relation is the same for Ernest and Cecily, the truth of iii and iv will follow from the truth of i and ii.

[17]Strawson writes "the aim, if not the achievement, of securing uptake is essentially *a standard, if not an invariable*, element in the performance of the illocutionary act" (1964: 449). Much later, Szabó and Thomason simply assert "Austin may or may not be right about warnings, but assertions, promises, and commands can certainly be made without the audience recognizing them" (2019: 218). Austin's claim has, however, been defended and applied in the philosophical literature on silencing: see, for example, Langton (1993), Hornsby (1994), Hornsby and Langton (1998).

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

For if Ernest can tell Cecily that *r* without her recognizing the force and content of the illocution, then Ernest can know that he has told Cecily that *r* and yet not be in a position to know that Cecily knows she has been told that *r*. In that case, we would not be able to establish SCK-2. Thus, I claim that if Austin is wrong, then we will not be able to establish SCK-2.

This argument might seem too fast. One might object: suppose that Austin is wrong, and Ernest can successfully tell Cecily that *r* without uptake being secured. Let us call such an action a telling′. And now suppose that in our description of A, we simply add that uptake has successfully occurred. Thus, we set:

(A–4) Ernest and Cecily have been talking. Ernest tells′ Cecily that he intends to return to the same place at the same time tomorrow to continue the conversation.
(C–4) Cecily has heard and understood what Ernest told′ her.

Surely, (the objection proceeds) if A–3 can serve as a basis for Strong Common Knowledge if Austin is right, then A–4 can serve as a basis for Strong Common Knowledge if Austin is wrong. In other words, if we can have Strong Common Knowledge if Austin is right, then we can have Strong Common Knowledge even if Austin is wrong. Remarkably, this is not so. To show that A–4 satisfies SCK-2, we'd need to establish:

i′ If Ernest knows A–4, then Ernest thereby knows that Cecily knows A–4.
ii′ If Cecily knows A–4, then Cecily thereby knows that Ernest knows A–4.

I take it that i′ is true. ii′, however, fails. The problem is that even if Cecily knows that she has heard and understood what Ernest has told′ her, she does not thereby know that Ernest knows that she has heard and understood. So, we do not have SCK-2 in the case of A–4.[18] I conclude that Austin* must be correct for A–3 to serve as a basis for Strong Common Knowledge.

I've indicated that we can modify Lewis' example so as to satisfy clause SCK-3. Suppose we set:

(A–6) Ernest and Cecily have been talking. Ernest promises Cecily that he will return to the same place tomorrow with his pet lizard.
(C–6) Ernest has promised to bring a reptile to the same place tomorrow.

Promises are illocutionary acts. If we adopt Austin*, then the argument of this section will show that A–6 satisfies clauses SCK-1 and SCK-2 of the definition of Strong Common Knowledge. What about clause SCK-3? This requires that it is true of both Ernest and Cecily that if they know A–6, they thereby know C–6. Recall that we are allowed to assume that Ernest and Cecily both have relatively ordinary background

[18]This might suggest yet a further modification:

(A–5) Ernest and Cecily have been talking. Ernest tells′ Cecily that he intends to return to the same place at the same time tomorrow to continue the conversation. Cecily has heard and understood what Ernest told′ her. Ernest knows that Cecily has heard and understood what he told′ her.

But now the problem is that I′′ If Ernest knows A–5, then Ernest thereby knows that Cecily knows A–5 will not be true.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

knowledge. Thus, we assume that each of them knows that lizards are reptiles. Thus, it will be true of each of them that in knowing A-6, they thereby know C–6.[19] So, assuming Austin*, A–6/C–6 is an example of Strong Common Knowledge.

My claim in this Section has been twofold. First, I've argued that if we assume Austin*, then A–3 satisfies SCK-1 and SCK-2, and A–6/C–6 satisfy SCK-1, SCK-2, and SCK-3. Second, I've argued that if we reject Austin*, neither A–3 nor A–6 can be a basis for Strong Common Knowledge. Austin's conception of speech acts seems tailor-made to satisfy the "thereby" clause of SCK-2.

5. Illocutions addressed to larger audiences

I've just argued that if we accept a particularly strong version of Austin's claim that the performance of an illocutionary act involves the securing of uptake, then we have a class of situations which can count as bases for Strong Common Knowledge: situations in which *two* people are talking, and one, addressing the other, intentionally performs some illocutionary act. However, the situation is more complicated if we seek a basis for Strong Common Knowledge in a *group* consisting of more than two.

(A–7) Ernest, Gwendolen, and Cecily have been talking. Ernest promises them that he will return to the same place tomorrow with his pet lizard.
(C–7) Ernest has promised to bring a reptile to the same place tomorrow.

To see why an audience of more than one person complicates the situation, recall that in order to establish SCK-2, we'll need to establish that:

(3) It is true of every person X in G (where G is a set whose members are, in this case, Ernest, Gwendolen, and Cecily) that: if X knows that A–7 holds, then X will thereby know that everybody in G knows that A–7 holds.

For example, we'll need to be sure, among other things, that it is true that:

(4) If Gwendolen knows A–7, then Gwendolen thereby knows that Cecily knows A–7.

I propose to take for granted that the first sentence of A–7 ("Ernest, Gwendolen, and Cecily have been talking") satisfies SCK-2. The difficulty lies in establishing the component of SCK-2, which concerns Ernest's speech act.

To see why there is a difficulty here, we must ask what is involved in what Austin describes as "bringing about the understanding of the meaning and of the force of the locution" in this case. What, exactly, does Gwendolen know when she knows that Ernest has "promised them to φ"? We might try the following analysis:

[19]One might object: knowing A–6 puts Cecily in a position to infer C–6, but what if she doesn't happen to take time to consciously consider this particular inference? As I indicated at the end of Section 3, I'm using 'know' in such a way that it can be true to say that Gwendolen knows that there are not exactly 729 unicorns, even if Gwendolen has never paid attention to this specific statement. I believe the same phenomenon is at work here: anybody competent with the relevant concepts who knows that this pet is a lizard thereby knows that this pet is a reptile, even if they have never asked themselves whether that particular statement is true.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

CONJUNCTIVE ANALYSIS 1 (CA1) In situation A–7 (where Ernest is talking to both Gwendolen and Cecily), when Gwendolen knows that Ernest has promised them to φ she knows that:

a. Ernest has brought about in Gwendolen an understanding that he has made a promise to φ, and
b. Ernest has brought about in Cecily an understanding that he has made a promise to φ.

In effect, CA1 assumes that Austin is correct in thinking that for a promise to be made, uptake must be secured, and then treats "Ernest has promised Gwendolen and Cecily to φ" as a simple conjunction of "Ernest has promised Gwendolen to φ" and "Ernest has promised Cecily to φ."

But although CA1 is already a stronger account of the speech act of promising than Austin's opponents would be prepared to countenance, CA1 is both unlikely as an interpretation of what is involved in the speech act of promising specified in A–7, and it is also insufficient to render A–7 a basis for common knowledge. The problem is that CA1 does not ensure that Gwendolen knows that Cecily knows that Ernest has made the promise in question, nor vice versa. And this means that (4) will not be satisfied. For in knowing A–7 so understood, Gwendolen will know that Cecily knows that a promise has been made to her (Cecily), but Gwendolen will not know whether Cecily knows that a promise has been made to Gwendolen. So, Gwendolen won't know that Cecily knows A–7.

Not only does CA1 prevent us from establishing (4), but CA1 also strikes me as an inadequate analysis of what is involved in the speech act described in A–7. Intuitively, when Ernest makes a promise to Gwendolen and Cecily, Gwendolen and Cecily will each know that the other is aware of the promise (the promise is, at least paradigmatically, public). Ordinarily, this is part of what it is to promise *them* that he will φ.

I propose that we can make A–7 satisfy SCK-2 if we insist that, in promising Gwendolen and Cecily that he will φ, Ernest brings it about that both of them know that he has promised *them* to φ.

If this condition is satisfied, then if Gwendolen knows A–7, she will thereby know that Cecily knows A–7: for if Gwendolen knows that Ernest has promised them that he will φ, then she thereby knows that Cecily knows that he has promised *them* that he will φ, which is precisely what is required for SCK-2.

I've italicized "them" to emphasize the role that this word is playing. The idea, in effect, is that in order to secure uptake for certain speech acts directed at an audience of more than one person, it is necessary not simply that the audience grasp the force and the content of the utterance, but that *they* recognize *themselves* as the audience for the utterance. The act of recognition is essentially plural. We cannot represent what the audience grasps with any finite conjunction of propositions without making use of plural pronouns.

One might think my insistence on such irreducibly plural knowledge is too hasty, and that we can ensure that SCK-2 is satisfied by generating a more complicated conjunction of facts that do not essentially involve plural pronouns.

One might, for example, propose the following:

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

CONJUNCTIVE ANALYSIS 2 (CA2) In situation A–7 (where Ernest is talking to both Gwendolen and Cecily), when Gwendolen knows that Ernest has promised them to φ she knows that:

a. Ernest has brought about in Gwendolen an understanding that he has made a promise to her to φ;
b. Ernest has brought about in Gwendolen an understanding that he has made a promise to Cecily to φ;
c. Ernest has brought about in Cecily an understanding that he has made a promise to her to φ;
d. Ernest has brought about in Cecily an understanding that he has made a promise to Gwendolen to φ.

So, now let us suppose that Gwendolen knows A–7, and so she knows all four clauses of CA-2. Will this be enough to ensure the satisfaction of (4)? In other words, does Gwendolen thereby know that Cecily knows all four of these? The answer to this is "No," for suppose that Gwendolen (who, we are supposing, knows A–7), in trying to determine whether Cecily knows A–7, asks herself: "Does Cecily know clause b of CA-2?" But as far as Gwendolen knows, that which Cecily knows is recorded in clauses c and d of CA-2, and neither of these puts Gwendolen in a position to know that Cecily knows that Gwendolen knows that Ernest has made a promise to Cecily.

I've claimed that we cannot give an account of what is involved in the speech act described in A–7, which will satisfy SCK-2 by specifying a *finite* conjunction of facts that do not make essential use of plural pronouns. Observe, however, that we can do so with an infinite conjunction:

INFINITE CONJUNCTION In situation A–7 (where Ernest is talking to both Gwendolen and Cecily), when Gwendolen knows that Ernest has promised them to φ she knows that:

a. Ernest has promised Gwendolen that φ and Ernest has promised Cecily that φ;
b. Each knows that Ernest has promised Gwendolen that φ and he has promised Cecily that φ;
c. Each knows that each knows that Ernest has promised Gwendolen that φ and he has promised Cecily that φ;
d. . . . And so on.

I've claimed that the proper extension of Austin's notion of uptake to an audience of more than one person requires us to maintain that Ernest will have failed to promise Gwendolen and Cecily that φ unless he brings it about that Gwendolen and Cecily know that he has promised *them* to φ, where this *them* cannot be finitely analyzed. So understood, A–7 can serve as a basis for common knowledge, properly so-called, that C–7.

This claim extends to other illocutionary acts directed at more than one person. Suppose that Ernest is attempting to warn Cecily and Gwendolen of some imminent threat (we suppose the three of them are within earshot of each other). It is surely possible that he could succeed in warning one but fail to warn the other. Furthermore, I think it is also possible that he could succeed in warning each of them, but not bring it

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

about that each knows that the other has been warned (perhaps he employs a gesture or phrase known to each, but which each doubts the other knows). In this case, the correct thing to say is that although he has warned both of them, he has failed to warn *them*. And this difference matters. Suppose that if Ernest succeeds in warning Cecily and Gwendolen of the impending danger in such a way that the warning becomes common knowledge in the manner I've been describing, then they will have time to turn and flee. Whereas (to continue the hypothetical) if Gwendolen recognizes the warning but is in some doubt about whether Cecily has recognized the warning, then Gwendolen will pause to try to rescue Cecily from danger, a pause which, we suppose, makes the difference between life and death. In this case, although Ernest managed to warn them several times, he did not succeed in warning *them*, with fatal consequences.[20]

Not every case in which one performs an illocution with an audience of more than one person will fit this pattern, for it is crucial that the target of the illocutionary act be the audience considered collectively. Suppose, for example, that Ernest, Cecily, and Gwendolen have been talking. Ernest (with Cecily still present) orders Gwendolen to φ. It is perfectly plausible to imagine a situation in which Ernest successfully orders Gwendolen to φ without Cecily knowing that he has done so (Ernest might, for example, disguise the order by using a special code known only to himself and Gwendolen, or Cecily might take it as a request, not noticing the tone of voice which Ernest uses). But in this case, Ernest made no attempt to order *them* to do anything, so there is no counterexample to my claim.

However, in a case in which Ernest is talking to Cecily and Gwendolen and orders *them* to φ, I need it to be the case that he will have failed to order them to φ if each takes themselves to have been ordered, but is in doubt about whether the other has been so ordered. And there is indeed a difference here. Suppose Ernest attempts to order them to load a hundred crates of such-and-such into the truck. Each mistakenly thinks this order has been issued to them severally, so that night each loads a hundred crates. And now, two hundred crates leave the warehouse, instead of the one hundred Ernest had intended. I think it is correct in this case to say that although he somehow managed to issue an order to each of them, he failed to issue the intended order to *them*.

6. Does common knowledge ever come about in this way?

In Section 3, I modified Lewis' original definition of common knowledge so as to define what I called Strong Common Knowledge. The satisfaction of that set of conditions would entail the satisfaction of the CK hierarchy. In Section 4, I observed that cases which are often taken to be paradigmatic for common knowledge, in which some everyday object is in the public view of a group of ordinarily equipped perceivers, do not appear to meet these conditions. I suggested a different example (§4–§5): I've claimed that, on a strong interpretation of Austin's claim (Austin*) that the performance of an illocution involves securing uptake, then common knowledge can be achieved in the course of a conversational exchange. I've also claimed that if Austin* is wrong, then

[20]It seems to me that if one wanted to do it, it would be quite hard to warn each of Gwendolen and Cecily about something, but not warn *them* of it. It might be slightly easier to warn just one, while leaving the other ignorant of the danger; nevertheless, this would still be harder than warning them. This seems to me to constitute some evidence not only that the distinction between warning each and warning them is real, but that the default case is that such publicly performed speech acts are addressed to a plural audience, and their successful performance depends on their being recognized as such.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

Strong Common Knowledge cannot come about as the result of a conversational exchange.

I have not, so far, argued that the conditions which define Strong Common Knowledge can ever, in fact, be satisfied. For I have not given an argument defending Austin*. If you are one of the many who reject Austin's claim, then my argument in fact provides reason to be sceptical that common knowledge, properly so-called, could ever come about.

I do not have a direct argument to make in defense of Austin*. However, in this Section, I draw attention to the striking fact that one of the so-called paradoxes of common knowledge bears directly on the issue of illocutionary acts and common knowledge.

According to the "Coordinated Attack Problem," we are to imagine two army divisions encamped on either side of an enemy.[21] If both divisions attack at the same time, the enemy will be vanquished. If only one division attacks, that division will suffer a catastrophic defeat. So, the general of each division will attack only if he knows that the other will do so. Unfortunately, the generals can communicate only via a messenger, who must sneak through enemy lines. Thus, each time a message is sent, there is a nonzero chance that the messenger is intercepted. We are to suppose that the generals believe the messenger.

To get a feeling for the problem, suppose that General Moncrieff sends a message to General Cardew outlining his plans to attack.[22] Suppose this message makes it through enemy lines and Cardew believes it. Now, we've assumed that each general will attack only if he knows that the other will. And since Moncrieff does not know whether Cardew has received his message, he does not know whether Cardew will attack. So, Moncrieff will not attack. So, suppose that Cardew writes back to confirm that she has received the message and intends to attack. Suppose that this message gets through and Moncrieff believes it. But now, Cardew will not attack. For she knows that Moncrieff will not attack unless Moncrieff knows that she has received his message and intends to attack, and (since there is a chance her message was lost) she does not know that Moncrieff knows that she plans to attack. Recognizing this, Moncrieff sends a message back to Cardew confirming that her message has been received. But now it won't be rational for Moncrieff to attack unless he has confirmation that Cardew has received his message, for he knows that Cardew won't attack unless she has received his message, and since there is some chance his message has been lost, he won't attack unless he receives confirmation from her.

Unsurprisingly, this pattern of reasoning iterates. It can be shown formally[23] that: the generals will attack only if they commonly know that each plans to attack, and that: no matter how many messages are sent, the generals cannot achieve such common knowledge. It follows that no matter how many messages they exchange, the generals will never be in a position to attack the enemy. The "paradox" arises from the clash between this result and the intuitive thought that surely, after messengers have gone back and forth between the generals a few times, each general can be confident about the intentions of the other, and it can be perfectly sensible for the generals to attack.

[21]My presentation of the problem draws mainly on Lederman (2018a), although he presents the problem in terms of common belief, not common knowledge.

[22]I emphasize that the discussion in this paragraph is informal.

[23]The language of this paragraph is essentially that of Lederman (2018a: 3).

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

To see how the coordinated attack problem bears on the question of the nature of speech acts, consider what difference it makes that the generals are on opposite sides of the enemy and able to communicate only via messenger. Suppose instead that the generals are in the same room when they discuss their plans. Once their plans are hatched, each will depart for their respective divisions to prepare to attack (or not). We continue to assume that neither will attack unless they know that the other intends to attack, from which it will again follow that the generals will attack only if it is common knowledge that they intend to do so.

In the original version of the problem, in which the generals are situated on opposite sides of the enemy, common knowledge cannot be achieved. I've argued that if we accept Austin*, then it is possible for two or more people engaged in a face-to-face conversation to achieve common knowledge. If that argument is correct, then it follows that if Austin* is a correct analysis of illocutionary acts and such illocutionary acts are possible, then when the generals are in the same room, there will be no paradox.[24] That is, if illocutionary acts, as described by Austin*, are possible, then the face-to-face generals can rationally act in a coordinated fashion on the basis of their common knowledge of a plan of attack.[25]

However, if Austin is wrong, then the paradox remains even if the generals are in the same room. For if Austin is wrong to claim that the performance of an illocutionary act requires the securing of uptake, then it will be possible for Ernest to "tell" Cecily that *r*, without it being the case that she knows that she has been "told" *r* by Ernest. It is precisely this possibility that is modeled, in the coordinated attack problem, by the possibility that the messenger does not make it through enemy lines. Just as Moncrieff can *send* a message to Cardew without being sure that Cardew has received it, so too (if Austin is wrong) can Ernest *"tell"* Cecily that *r* without being assured that Cecily has heard and understood what she has been "told," even when Ernest and Cecily are face-to-face. In other words, if Austin is wrong, then there is always a small chance that any "message," which an agent has "sent" in the course of a face-to-face conversational exchange, will fail to be "received." So, if Austin is wrong, then even if Ernest and Cecily are in the same room, in order for it to become common knowledge that Ernest has "told" Cecily that *r*, he'll need some indication from her that she knows that he has "told" her that *r*. But confirming that one has understood what one has been "told" – even when one is in the same room – is itself a speech act. And since (on the assumption that Austin is wrong) it is possible to perform such a speech act without uptake being secured, then in order for Cecily to be sure that Ernest has understood her confirmatory speech act, it'll have to be the case that Ernest confirms to Cecily that he has understood her confirmatory speech act. And so on indefinitely.[26]

The current claim is, therefore, that if Austin is wrong, then the plight of any two interlocutors, even two who are in the same room, is exactly the same as the plight of the two generals on opposite sides of the opposing army. At best, Cecily can know what it is

[24]Note that we do indeed need Austin* here. For one might think that if the performance of an illocutionary act requires uptake, then it would follow that such an act has been successfully performed once Moncrieff's initial message is received by Cardew. But this is not so if, as I urged, we understand Austin's claim to be that articulated in Austin* (Section 4). For in that case, 'the securing of uptake' requires that Moncreiff know that he has secured uptake for his illocutionary act.

[25]The situation would, I believe, be roughly the same if we added further generals. Three generals encamped on three sides of an army would, obviously, be incapable of elevating any plan to the status of common knowledge. But if the argument of Section 5 is correct, the three generals would be able to come to commonly know some plan of attack if they were able to meet face-to-face-to-face.

[26]Thanks to an anonymous referee for encouraging me to make this point more clearly.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

that Ernest is attempting to "tell" her, Ernest can know that Cecily knows this, Cecily can know that Ernest knows that Cecily knows, and so on, but we can continue the series only for as long as their patience at exchanging confirmatory remarks holds out.[27] If Austin is wrong, one can successfully perform a speech act but still be in some doubt about whether one's audience has recognized the force and content of that speech act. In that case, as far as achieving common knowledge is concerned, the plight of any two conversationalists is the same as the plight of the two generals on opposite sides of the enemy.[28]

My view is, therefore, that the coordinated attack problem in fact gives us some reason in favor of thinking that Austin's account of illocution is correct. If all communicative exchanges must be modeled on the case of the two generals encamped on opposite sides of the army, then common knowledge can never be achieved through communication. But the coordinated attack problem then forces us to confront the question: How can it ever be rational to engage in coordinated action if common knowledge can never be achieved?[29] An alternative becomes available once we recognize that the plight of the two generals is not the correct way to model all communicative exchanges. I've argued that genuine common knowledge can be achieved in the course of a conversation if we accept Austin*. This doesn't itself provide a solution to the problem confronting the generals encamped on opposite sides of the enemy. The point is that, on the one hand, if we accept Austin*, interlocutors can generate common knowledge in the course of an in-person conversation and thereby rationally engage in coordinated action. On the other hand, if Austin is wrong about uptake, then genuine common knowledge cannot be achieved, and any communicative exchange, even one between people in the same room, must resemble that between the generals in the coordinated attack problem, continually sending one-way messages to each other and never able to secure common knowledge. The most straightforward path through these claims is to accept that sometimes coordinated action is rational only if genuine common knowledge can come

[27]I am now making a stronger claim than the one I made in Section 4. There I argued that if Austin is wrong about uptake, then the modified Lewisian conditions I've called Strong Common Knowledge will not be satisfied. I'm now making the slightly stronger claim that if Austin is wrong, it is unclear how common knowledge, properly so-called, could ever arise in the course of a conversation.

[28]This is indeed how the coordinated attack problem is sometimes interpreted. For example: "When Alice suggests meeting at noon in the coffee shop and Bob says 'OK,' this is usually enough to ensure that Alice and Bob will succeed in meeting the next day. But their agreement is not common knowledge between them because Bob did not get his confirmation confirmed" (Binmore and Samuelson 2001: 8). In her discussion of this problem, Jankovic notes that the coordinated attack problem is "often described as a case in which agents who are able to communicate cannot achieve common knowledge of the communicated fact" (Jankovic 2014: 493–4). Thus, she quotes, for example, van der Grijn: "Making use of epistemic logic, it is provable that whenever the smallest uncertainty of message delivery is present, common knowledge via communication is impossible" (Grijn 1994: 1). Jankovic herself addresses the question of how communication itself could be possible if common knowledge cannot be secured. Chant and Ernst (2008) agree that the coordinated attack problem shows that common knowledge cannot be achieved through communication, and so ask what epistemic conditions, short of common knowledge, might be required for collective action to be rational. I agree that there might be situations in which coordinated action is rational in the absence of common knowledge. My claim here is simply that if Austin is right, common knowledge can be achieved through face-to-face communication in situations where such knowledge is required for coordinated action.

[29]For example, Chant and Ernst (2008) attempt to show how coordinated action could be rational in the absence of common knowledge, Jankovic (2014) attempts to show how communication could be possible if it is impossible to achieve common knowledge, and Lederman (2018a) argues that coordinated action can be perfectly rational in the absence of both common knowledge and common belief.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

about, but to hold that in-person communication, deploying speech acts as described by Austin*, ensures that common knowledge can, at least in some circumstances, arise.[30]

7. Conclusion

The question I've addressed in this paper is whether a group of individuals could ever come to have common knowledge, properly so-called. I started with Lewis' account, for it is one of the few that would allow us to see how a group of agents could achieve (what he calls) common knowledge. By modifying his account, I generated a set of conditions the satisfaction of which would entail that a group of agents has genuine common knowledge. It turns out to be rather hard to satisfy these conditions. Situations in which an ordinary object is in plain view of a group of perceivers do not furnish us with straightforward examples. I argued instead that on a suitably strengthened version of Austin's claim that the performance of an illocutionary act requires the securing of uptake, the modified Lewisian conditions can be satisfied in the case of a conversational exchange between two people. However, extending Austin's notion to the case of a conversation involving more than two people is even more demanding. We must assume that an audience consisting of more than one person is capable of an essentially plural act of recognition: they can recognize *themselves* as the audience of a speech act. I closed by suggesting that the coordinated attack problem, in fact, gives some reasons in favor of Austin's account of illocutions. According to one interpretation of the coordinated attack problem, common knowledge cannot come about in the course of communication if our picture of communication is dictated by the image of two generals on distinct sides of an army sending a messenger, who could be intercepted at any moment, back and forth between one another. Austin's conception of illocution, in which speech acts require uptake, provides an alternative model in which common knowledge can indeed come about in the course of a conversation.

[30] At the core of this line of thought is the following argument. The rationality of coordinated action depends on the possibility of common knowledge. Generals who can meet face-to-face can achieve common knowledge only if illocutionary acts, as described by Austin*, are possible. Since it can manifestly be rational for generals who are able to meet face-to-face to launch a coordinated attack, it follows that illocutionary acts, as described by Austin*, must be possible. There is, however, at least one complication with this way of presenting the argument. Like others, I believe that it can be rational for the generals to attack in the original version of the coordinated attack problem, for like others, I believe we have to accept that it can sometimes be rational for the generals to attack even if they cannot achieve common knowledge of their plans. My own preferred solution to the coordinated attack problem insists that our account of what can make it rational for the generals to attack in the original version of the coordinated attack problem must presuppose the rationality required for the kind of genuine common knowledge which, I have claimed, is available on the basis of Austinian illocutionary acts. Specifically, I claim that after a certain number of messages have been exchanged between interlocutors who are not face-to-face, it can be rational to proceed *as if* an illocutionary act has been performed. Thus, if I propose, by text, to meet you for coffee tomorrow at the usual place at 2 pm, and you respond with an upraised thumb, it is rational to proceed *as if* an illocutionary exchange had taken place, even though, strictly speaking, common knowledge has not been established. If the stakes are higher, it might take more messages (or more trust) for it to be rational to proceed as if an illocutionary exchange had taken place. But this can only be rational, I would argue, on the assumption that we can, and do, very often achieve genuine common knowledge in the course of face-to-face conversational exchanges. This raises the question: under what kinds of circumstances is it possible for genuine illocutionary acts to be performed? I would argue (although I cannot defend the view here) that what matters for the securing of uptake is that the conversational participants are in direct contact with one another: so it would be possible for full-blooded illocutions to be performed via video conference or phone calls, but not via text, or email, or letter, or even a trusted intermediary. Communication in those cases requires us to proceed as if illocutions have taken place.

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press

If my argument is correct, then it shows that it is not easy to find straightforward examples of common knowledge, properly so-called. I've attempted to show that nothing short of Austin's account of illocution, suitably interpreted, can secure common knowledge in the course of a conversation. I tend to think this speaks in favor of Austin's account. However, I emphasize that another perfectly reasonable interpretation of my argument would be to reject both the idea that common knowledge can ever come about and to reject my interpretation of Austin's conception of uptake.

Acknowledgements. I'd like to thank audiences at the Alabama Philosophical Society and the Joint Session of the Aristotelian Society and the Mind Association for listening to earlier versions of this material. I'd also like to thank the referee and area editor for this journal for their generous and unusually helpful comments. Finally, I thank Jennifer Lockhart and Arata Hamawaki for the many conversations which inspired and improved this paper.

References

Aumann R. (1976). 'Agreeing to Disagree.' *Annals of Statistics* **4**, 1236–9.
Austin J.L. (1962). *How To Do Things With Words*. Cambridge, MA.: Harvard University Press.
Barwise J. (1988). 'Three Views of Common Knowledge.' In: Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge. Ed. by M Y Vardi. San Francisco: Morgan Kaufman.
Binmore K. and Branderburger A. (1988). *Common Knowedge and Game Theory*. ST/ICERD Discussion Paper 88/167. London School of Economics.
Binmore K. and Samuelson L. (2001). 'Coordinated Action in the Electronic Mail Game.' *Games and Economic Behavior* **35**, 6–30.
Chant S. and Ernst Z. (2008). 'Epistemic Conditions for Collective Action.' *Mind* **117**, 549–73.
Cubitt R.P. and Sugden R. (2003). 'Common Knowledge, Salience and Convention: A Reconstruction of David Lewis' Game Theory.' *Economics and Philosophy* **19**, 175–210.
Fitzgerald F.S. (2025). *The Great Gatsby*. Oxford World's Classics.
Gilbert M. (1989). *On Social Facts*. Princeton, NJ: Princeton University Press.
Grijn F van der (1994). (Im)possibility of a coordinated attack. ILLC Publications: Preprints and Other Reports. Universitet van Amsterdam.
Hornsby J. (1994). 'Illocution and Its Significance.' In S. L. Tsohatzidis (eds), *Foundations of Speech Act Theory: Philosophical and Linguistic Perspectives*, pp. 187–207. New York: Routledge.
Hornsby J. and Langton R. (1998). 'Free Speech and Illocution.'*Legal Theory* **4**, 21–37.
Jankovic M. (2014). 'Communication and Shared Information.' *Philosophical Studies* **169**, 489–508.
Langton R. (1993). 'Speech Acts and Unspeakable Acts.' *Philosophy and Public Affairs* **22**(4), 293–330.
Lederman, H. (2018a). 'Two Paradoxes of Common Knowledge: Coordinated Attack and Electronic Mail.' *Noûs* **52**, 921–45.
Lederman H. (2018b). 'Uncommon Knowledge.' *Mind* **127**, 1069–105.
Lewis D. (1969). *Convention*. Cambridge, MA.: Harvard University Press.
Monderer D. and Samet D. (1989). 'Approximating Common Knowledge with Common Beliefs.' *Games and Economic Behavior* **1**, 170–90.
Strawson P.F. (1964). 'Intention and Convention in Speech Acts.' *The Philosophical Review* **73**(4), 439–60.
Szabó Z.G. and Thomason R.H. (2019). *Philosophy of Language. Cambridge Textbooks in Linguistics*. Cambridge: Cambridge University Press.
Vanderschraaf P. and Sillari G. (2023). 'Common Knowledge.' InN.Z. Edward and U. Nodelman (eds), *The Stanford Encyclopedia of Philosophy*. Winter 2023. Metaphysics Research Lab, Stanford University.

Thomas Lockhart is an associate professor in philosophy at Auburn University. His research interests include epistemology, especially epistemological disjunctivism, and Frege. Email: lockhartt@auburn.edu

Cite this article: Lockhart T. (2026). "Achieving Common Knowledge." *Episteme* **23**, 173–194. https://doi.org/10.1017/epi.2026.10101

https://doi.org/10.1017/epi.2026.10101 Published online by Cambridge University Press