

*Practical Moore Sentences**MATTHEW MANDELKERN
All Souls College, Oxford**Abstract**

I discuss what I call *practical Moore sentences*: sentences like ‘You must close your door, but I don’t know whether you will’, which combine an order together with an avowal of agnosticism about whether the order will be obeyed. I show that practical Moore sentences are generally infelicitous. But this infelicity is surprising: it seems like there should be nothing wrong with giving someone an order while acknowledging that you do not know whether it will be obeyed. I suggest that this infelicity points to a striking psychological fact, with potentially broad ramifications concerning the structure of norms of speech acts: namely, when giving an order, we must *act as if* we believe we will be obeyed.

1. Introduction

Moore (1942) observed that there is something peculiar about asserting a sentence like (1):

- (1) #It’s raining but I don’t believe it’s raining.

This is at least *prima facie* puzzling, because (1) can certainly be true. The same point goes for nearby variants on Moore’s sentence, like (2)–(4).

- (2) #It’s raining but I don’t know it’s raining.
 (3) #For all we know it’s not raining, but it is raining.
 (4) #It might not be raining but it is. (Wittgenstein 2001/1953: II.x.109)

The lesson of these *Moore sentences*, many have thought, is that you must bear the right cognitive attitude towards what you assert—you must know, or at least believe, what you assert. And given plausible assumptions about the logics of knowledge and belief, it is impossible to know or believe the content of a Moore sentence. Suppose, for instance, that you knew (2). Then, assuming knowledge distributes over conjuncts, you would know that you don’t know it’s raining. By the factivity of knowledge, it would follow, absurdly, that you know it’s raining and that you don’t know it’s raining. Similar considerations go for the other variants.¹

*Many thanks to two anonymous referees for this journal, to the audience at DEON 2018 in Utrecht, and to Sander Beckers, Agnes Callard, Fabrizio Cariani, Hasan Dindjer, Jeremy Goodman, Daniel Harris, Brendan de Kenessey, Dilip Ninan, Jonathan Phillips, Jake Quilty-Dunn, Daniel Rothschild, and Zoltán Szabó for very helpful comments and discussion.

In this paper I discuss a parallel to Moore sentences in the practical domain. *Practical Moore sentences* are sentences like (5), which conjoins an order with an avowal of ignorance about whether it will be obeyed:

(5) #You must close the door, but I don't know whether you will.

Building on observations from Silk 2015, 2018, I argue that sentences like this are in general infelicitous. I argue that this infelicity is surprising, because there are many circumstances in which both conjuncts of a sentence like (5) seem to be true. Indeed, the situation is even more puzzling than for Moore sentences, because (unlike for Moore sentences) there are many situations in which both conjuncts of (5) are not only true but also knowable and believable; and, indeed, many situations in which each conjunct on its own is assertable, despite the fact that there are no situations where the whole conjunction is assertable. I argue that, like Moore sentences, practical Moore sentences teach us something new about attitudes: in giving an order, you must act towards your addressee as if you believe your order will be obeyed (whether or not you actually believe it will be). If this is right, it is a surprising fact about ordering, which may also have broad ramifications concerning the norms governing speech acts in general.

2. Practical Moore Sentences

Suppose Sue is ordering her student Liz to turn in her final paper by the end of the exam period. She might do so by saying one of the following:

- (6) You must turn in your final paper by the end of the exam period.
- (7) Turn in your final paper by the end of the exam period.

Or, in a peremptory mood, (8) or (9):

- (8) I order you to turn in your final paper by the end of the exam period.
- (9) I command you to turn in your final paper by the end of the exam period.

Now, Sue is an experienced teacher. She knows that students don't always get their papers in on time. She knows that Liz in particular is a chronic procrastinator. She wants Liz to know that she knows this, too, so that Liz understands the potential consequences. She might communicate this with any of the following:

- (10) You might not turn your paper in on time.
- (11) I don't know if you will turn your paper in on time.
- (12) For all I know, you won't turn your paper in on time.

She might carry on 'In that case, you'll fail the class'.

This is all perfectly quotidian. The observation at the heart of this paper is that there is something wrong with asserting the *conjunction* of any of the sentences in (6)–(9) together with any of the sentences in (10)–(12). Here is a sampling of possible combinations:²

- (13) #You must turn in your final paper by the end of the exam period, but you might not turn your paper in by then.
- (14) #You must turn in your final paper by the end of the exam period, but I don't know if you will turn your paper in by then.
- (15) #You must turn in your final paper by the end of the exam period, but for all I know you won't turn your paper in by then.
- (16) #You might not turn your final paper by the end of the exam period, but turn in your final paper by then.
- (17) #I don't know if you will turn in your final paper by the end of the exam period, but turn in your final paper by then.
- (18) #I order you to turn in your final paper by the end of the exam period, but you might not turn it in by then.
- (19) #I order you to turn in your final paper by the end of the exam period, but I don't know if you will turn it in by then.

Note that the infelicity does not depend on the order of the conjuncts: thus e.g. (20) is just as infelicitous as (13), and (21) is as bad as (19):

- (20) #You might not turn your paper in by the end of the exam period, but you must turn it in by then.
- (21) #I don't know if you will turn in your final paper by the end of the exam period, but I order you to do so.

The observation that conjoining deontic modals with corresponding avowals of ignorance leads to infelicity is due to Silk 2015, 2018.³ The range of present data suggests that that infelicity is just one part of a broader pattern. What is the generalization that these data point towards?

Sentences (6)–(9) are all order-giving sentences: sentences in which the speaker orders the addressee to do some action φ .⁴ Let us use 'ORDER(φ)' as a schema which ranges over uses of sentences which are orders to do φ :⁵ this includes sentences containing strong deontic necessity modals which are intended as orders ('You must do φ ', 'You have to do φ ', 'You can't do φ '; though not when these are restricted in a way which makes it clear that they are not intended to be orders, as in 'According to certain morally bad laws, you must kill your grandmother'); imperatives which are intended as orders ('Do φ ', though not as in permissive uses, as in 'Close the door! Leave it open! I don't care!'); and performative orders ('I order you to do φ ').

Sentences (10)–(12) are all proposals to leave it open that some state of affairs obtains. Let us use 'OPEN(φ)' as a schema which ranges over uses of sentences which propose to leave it open that φ is true: these include epistemic possibility claims ('Might φ ', 'Could φ '), as well as avowals of ignorance ('I don't know whether φ ', 'For all I know, φ ').

With this in hand, I propose that the right generalization to draw from the infelicity of sentences like (13)–(21) is the following:

Speech acts of the form $\ulcorner \text{ORDER}(\varphi) \wedge \text{OPEN}(\neg\varphi) \urcorner$ and of the form $\ulcorner \text{OPEN}(\neg\varphi) \wedge \text{ORDER}(\varphi) \urcorner$ are typically infelicitous.

This generalization explains the infelicity of (13)–(21). Moreover, it makes the right predictions about further data in the neighborhood. First, note that further permutations which have the form of practical Moore sentences are all, as far as I can tell, infelicitous. Thus to get one more case on the table, suppose Liz is storming out of Sue’s office in a huff. Sue wants Liz to close the door behind her, but knows that Liz might not listen to her. The following all sound quite weird in this situation:

- (22) #You might not close the door, but close it!
- (23) #You have to close the door after you, but you might not!
- (24) #You might not close the door, but I order you to do so!

And so on.

Second, nearby permutations on practical Moore sentences which do *not* amount to orders are *not* infelicitous, in line with our generalization. Consider what happens when we replace $\ulcorner \text{ORDER}(\varphi) \urcorner$ with a construction which tells one’s addressee what she *should* do, or what the speaker would *like* her to do, without rising to the level of an order, as in:

- (25) You should open the door.
- (26) You ought to open the door.
- (27) I want you to open the door.

I’ll schematize claims with this form ‘SHOULD(φ)’. Conjunctions of the form $\ulcorner \text{SHOULD}(\varphi) \wedge \text{OPEN}(\neg\varphi) \urcorner$ and of the form $\ulcorner \text{OPEN}(\neg\varphi) \wedge \text{SHOULD}(\varphi) \urcorner$ do not have the same infelicity as practical Moore sentences. So Sue might say any of the following to Liz:

- (28) You should close the door, but you might not.
- (29) You might not close the door, but you should.
- (30) I want you to close the door, but I don’t know whether you will.
- (31) You ought to close the door, but for all I know, you won’t.

And so on. These are all markedly more felicitous than practical Moore sentences; to make the contrast very clear, compare the following:

- (32) a. You should close the door, but I don’t know if you will.
b. You must close the door, but I don’t know if you will.

What is important for present purposes is that these sentences are structurally very much like practical Moore sentences, with the one striking difference that sentences of the form $\ulcorner \text{SHOULD}(\varphi) \urcorner$ clearly do not amount to *orders* to do φ (as evidenced by the felicity of conjunctions like ‘You should close the door, but you don’t have to’, or ‘I want you to close the door, but you don’t have to’).⁶ This provides confirmation,

again, that our generalization above is the correct one: the infelicity of (13)–(21) comes from the fact that they conjoin an *order* to φ with a proposal to leave it open that φ is false.

Further support for our generalization comes from contrasts between practical Moore sentences and similar sentences which contain strong deontic necessity claims and imperatives which do not count as orders. Among strong deontic necessity claims, we can distinguish *performative* from *descriptive* uses (see Lewis 1979, and Kaufmann 2012, §2.3.1 for recent discussion). Performative uses amount to orders; descriptive uses simply describe some system of norms, and thus will not always have the force of commands (in particular, when that system of norms does not bind the addressee). Descriptive uses are particularly easy to bring out with overt restrictors like ‘According to such and such body of law’, as in (33), but can also be brought out with the right kind of contextual set-up, as in (34), where it is clear that the lawyer is simply describing the existing system of laws, and not necessarily assuming it will be obeyed.

- (33) According to local custom, you have to take exactly two lumps of sugar in your coffee. But you should not feel bound by local custom, and for all I know you will take more than two!
- (34) a. [Client:] What is my legal obligation, and what do you expect me to do?
 b. [Lawyer:] You have to report your liability, but I don’t know if you will; you may prefer to push the limits of the law and just conceal it.

The present generalization predicts that *descriptive* uses of strong deontic necessity claims, unlike performative uses, will not necessarily be felt to be infelicitous when conjoined with the corresponding OPEN sentences. This seems to be correct: sentences like (33) and (34) (due to an anonymous referee for this journal) strike me as felicitous.

We find similar contrasts when we compare imperatives which are used to order with imperatives which are used instead to permit or acquiesce (see von Fintel & Iatridou 2017 for recent discussion of such uses). For instance, the two imperatives in (35) are intuitively not commands:⁷

- (35) Close the window! Don’t close the window! I don’t care at all.

If these were commands, the speaker would be commanding the addressee to both close the window and to leave it open, and hence would be commanding her to do something impossible—which, intuitively, she is not doing. Rather, the imperatives in (35) intuitively are something more like permissions to either close the window or leave it open. Since our generalization is about *orders*, not imperatives in particular, we predict that sentences like (35), with non-order giving uses of imperatives, can sound fine conjoined with the corresponding OPEN sentence. And, again, this seems to be correct:

- (36) Close the window! Don’t close the window! I don’t know whether you’ll close it or not.

(36) sounds fine. More generally, when we accompany an imperative with a clashing imperative, thereby making clear that neither is being used to give an order, conjoining the result with the corresponding OPEN sentence does not seem to lead to infelicity:

- (37) You might close the window; but close it, don't close it, what do I care?
- (38) You can have only one piece of fruit. Have a pear! Have an apple! I don't know which one you'll take.
- (39) You can have only one piece of fruit. Have a pear! Have an apple! You might have either one.

The felicity of these sentences suggests that the infelicity of the conjunctions we saw above, which conjoin imperatives with the corresponding OPEN sentences, does not stem from a feature of imperatives in general. It is not *generally* the case that conjunctions of imperatives with the corresponding OPEN sentences are infelicitous: that hypothesis would not explain the felicity of the sentences in (36)–(39), nor would it explain the infelicity of the conjunctions containing deontic modals and performative orders we saw above. Instead, the felicity of sentences like (36)–(39) provides further support for the generalization that there is something wrong with conjoining *orders* in particular with corresponding OPEN sentences. Imperatives provide one canonical way to give orders, but not all imperatives are used to give orders, and not all orders are given with imperatives; the infelicity of our sentences seems to track whether they are used to give orders, not whether they contain an imperative.⁸

It is important to note that, if we adopt the present generalization, we will need to think of the category of orders in a relatively broad way. For instance, to take an example suggested by an anonymous referee, consider (42):

- (42) #You might not come to my party, but come!

If we want to capture the infelicity of (42) with our generalization, we will need to count it as an order. A similar example, suggested by the same referee, comes from the variant on (19) in (43):

- (43) #I ask you to turn in your final paper by the end of the exam period, but I don't know if you will turn it in by then.

To account for the infelicity of (43) by way of our generalization, we need to count (43) as an order. It might be helpful for some taxonomic purposes to distinguish *requests* from *orders*, and to class (43) as a request, not an order. Likewise, it might be helpful to distinguish *invitations* from *orders* and to class (42) as an invitation, not an order. But the infelicity of (43) suggests that this is not the right level of taxonomic granularity for present purposes: rather, we should count performatives like 'I ask...' and invitations like those in (42) as orders, for the purposes of the present generalization. In general, I believe there are different speech act taxonomies which will be appropriate for different theoretical purposes: the present examples

suggest that the relevant level of taxonomy for the present generalization is relatively coarse grained, and counts things like (42) and (43) as orders.

It is worth noting, finally, that *conjunction* in particular does not play a key role in our generalization: consecutive assertions of $\lceil \text{ORDER}(\varphi) \rceil$, followed by $\lceil \text{OPEN}(\neg\varphi) \rceil$, or the reverse direction, are equally infelicitous:

- (44) #You must close the door. I don't know if you will.
 (45) #Close the door! You might not do so.

And so on. For simplicity, I will continue to focus on the single-sentence variants here, but everything I say applies to these sequential versions as well.

In sum, these contrasts support the claim that the source of infelicity in our core data is the fact that they are speech acts which combine an order to do some action with a proposal to leave it open whether that action will get done. Close variants which do not give orders, but instead give advice, describe a set of customs or laws, or give permissions are felicitous, just as our generalization predicts.

3. Why This Is So Weird

What is so weird about practical Moore sentences is that it seems like there are many situations in which practical Moore sentences can be true; in which they can be believed and known; and in which either conjunct, on its own, can be felicitously asserted—all this despite the fact that the conjunctions themselves are quite marked.

Taking each of these in turn: first, practical Moore sentences can be true (at least the variants with deontic modals, and plausibly also performatives—whereas imperatives are not obviously truth evaluable). In the situation described above, the following seems true:

- (46) Liz must turn in her final paper by the end of the exam period.

(46) seems like a true description of the requirements that bind Liz vis-à-vis her paper. And the following also seem true in the situation above:

- (47) Liz might not turn her paper in by then.
 (48) I [Sue] don't know if you will turn your paper in by then.

(47) is true because of facts about Liz: she is a procrastinator, chronically turning in papers late. We have no idea whether she will turn this one in on time; we can hope that she will, but we cannot be sure. (48) is true because of Sue's mental state: she simply doesn't know whether Liz will turn her paper in on time, given Liz's tendency to procrastinate. Assuming that truth of both conjuncts suffices for the truth of a conjunction, then, the conjunction of (46) with either (47) or (48) will be true in the situation described above. But these conjunctions will be practical Moore sentences. So a practical Moore sentence can be true.

Likewise, practical Moore sentences (crucially unlike Moore's own sentences) seem knowable and believable (setting aside, again, the imperative variants). Sue

knows her course requirements. So nothing stops her from knowing (46). And Sue knows what kind of a student Liz is: that Liz procrastinates, and might not turn her paper in on time. So nothing stops Sue from knowing both (47) and (48). Then, assuming that knowledge of both conjuncts suffices for knowledge of a conjunction, it follows that practical Moore sentences can be known (and thus believed). This makes their infelicity even more puzzling than that of Moore's paradox: not only are practical Moore sentences consistent (like Moore sentences); they are also (unlike Moore sentences) knowable (and thus believable), meaning that we cannot explain their infelicity in the same way that we account for the infelicity of Moore sentences.

Turning to assertability, things look even more puzzling. As we saw above, Sue is entirely within her rights to order Liz to turn her paper in on time: in the situation above, she has the right kind of authority, knowledge, etc. to tell Liz when she must turn in her paper by (despite knowing that Liz might disobey her). And she is also entirely within her rights to inform Liz that Liz might not turn her paper in on time: Sue knows this is so, and may have good reason to let Liz know that Sue knows it is so. So there are cases in which each conjunct of a practical Moore sentence is assertable. But there do not seem to be situations in which their conjunction is assertable.

4. What Practical Moore Sentences Are Not

This gives a sense of why practical Moore sentences are so puzzling. It's perfectly all right to give an order to do φ , even when you leave it open that (even after that order) your addressee may not do φ ; in the same situation, it's perfectly all right to assert that you leave it open that your addressee won't do φ ; but it does not seem permissible to do both of these things simultaneously.

Before offering a tentative account of the source of the infelicity of practical Moore sentences, I will discuss and dismiss a number of *prima facie* attractive, but, I argue, unsuccessful accounts.

4.1 Rhetoric

A natural thought about our conjunctions is that there is something rhetorically strange about them. After all, why highlight that you don't know whether your addressee will obey you right at the moment when you are telling them what to do? There is just something strange about putting these two speech acts together.

I think there is much to this thought, and a precisification of it will form the core of my account below. But we cannot help ourselves to this line in such a straightforward way. This is because the SHOULD variants on practical Moore sentences which we explored above are not infelicitous in the same way as practical Moore sentences. But $\lceil \text{SHOULD}(\varphi) \wedge \text{OPEN}(\neg\varphi) \rceil$, like a practical Moore sentence, puts together two thoughts which do not naturally go together. The fact that sentences with this form do not have the infelicity of practical Moore sentences shows that the infelicity of the latter cannot be simply due to the fact that they conjoin thoughts which somehow pull in opposing directions.

4.2 Mutual knowledge

A different natural thought about what's wrong with practical Moore sentences is that, though they can be known, they cannot be *mutually* known. It's natural to think that it is infelicitous to order someone to do φ when *they* know that *you* think they might not obey you. This would immediately account for the infelicity of practical Moore sentences: they can't be asserted, because they can't become mutually known.

But, again, things can't be this simple: it can be perfectly felicitous to order someone to do something even when it's clearly mutually known that they might disobey you. To illustrate the point, suppose that Mark is being kidnapped by some particularly vicious-seeming kidnappers. As he is being carried away over his kidnapper's back, he knows it is extraordinarily unlikely that they will let him go, whatever he does; and it is surely plausible that the kidnappers know that he knows this as well. Still, it seems perfectly felicitous for Mark to cry out:

- (49) Let me go!
 (50) You have to let me go right now!
 (51) I order you to release me!

So an order can be felicitous even when it is mutually known that the addressee might disobey. Even in such situations, however, the conjunction of the order with the claim that the addressee might disobey yields striking infelicity:

- (52) #You might not let me go, but let me go!
 (53) #You have to let me go right now, even though I don't know if you will!

4.3 The semantics of modals

A different thought would be that our phenomenon has something to do with the details of the semantics of modals. The idea would be that $\ulcorner \text{ORDER}(\varphi) \urcorner$ is, one way or another, semantically inconsistent with $\ulcorner \neg\varphi \urcorner$.⁹ Then $\ulcorner \text{ORDER}(\varphi) \urcorner$ and $\ulcorner \text{OPEN}(\neg\varphi) \urcorner$ would be felt to be in pragmatic tension, for the familiar Moorean reasons reviewed in the introduction. The most straightforward way of implementing this idea would be to simply assume that strong deontic necessity modals are associated with a reflexive accessibility relation when they are used to give an order, and to extend a parallel story to imperatives and performatives; then $\ulcorner \text{ORDER}(\varphi) \urcorner$ would entail φ and so would be inconsistent with $\ulcorner \neg\varphi \urcorner$.

This is an interesting possibility, and would be a striking and surprising conclusion. I do not think this response is right, though. I will start by giving two general arguments against this approach, and then explore in more detail two possible implementations of this idea.

The first argument is that, as we have seen, practical Moore sentences are not only consistent (like Moore sentences) but also intuitively believable and knowable. This suggests, again, that the situation here is unlike the situation in the case of Moore sentences, and in particular that $\ulcorner \text{ORDER}(\varphi) \urcorner$ is not semantically inconsistent

with $\ulcorner \neg\varphi \urcorner$; if it were, we would expect that $\ulcorner \text{ORDER}(\varphi) \wedge \text{OPEN}(\neg\varphi) \urcorner$ would be not only unassertable, but also unknowable.

A second, closely related argument against this approach is that, in ordering someone to φ , there is nothing wrong with *thinking* that they might not obey you, or with saying *to someone else* that they might not obey you. Suppose again that some notorious kidnappers are carrying you off. You shout: ‘Let me go!’ At the very same time, you can think or mutter *to yourself* that they might not let you go. Indeed, it seems positively irrational in this case *not* to have this thought, even though, as we have seen, it is very odd to say ‘Let me go, though you might not!’ Likewise, Sue might order Liz to turn her paper in on time—‘Liz, turn your paper in on time!’—and at the same time quietly turn to her teaching assistant and say: ‘I don’t know if she’s going to do it. Please let me know if she doesn’t.’ So it seems that what leads to infelicity is ordering someone to do φ and then acting *towards that person* as though she might not do φ . There does not seem to be anything wrong with ordering someone to do φ and then thinking to oneself, or saying to someone else, that they might not do it. But this is not what we would expect if $\ulcorner \text{ORDER}(\varphi) \urcorner$ were semantically inconsistent with $\ulcorner \neg\varphi \urcorner$. If that were so, then when you order someone to do φ , it should be felt to be incoherent to even think to yourself that φ might not happen, or to say this to someone other than your addressee. This is exactly what we find in the case of Moore’s paradox: if I sincerely tell Mark ‘It’s raining’, then it is incoherent for me to simultaneously think to myself ‘I don’t know that it’s raining’, or to say this to someone else. But this is not what we find in the case of practical Moore sentences, which suggests $\ulcorner \text{ORDER}(\varphi) \urcorner$ is not semantically inconsistent with $\ulcorner \neg\varphi \urcorner$.

With these general remarks in hand, let me explore two specific semantic approaches to accounting for practical Moore sentences that may look promising. The first notes an apparent similarity between practical Moore sentences and Wittgenstein’s variant on Moore’s sentence, which has the form $\ulcorner \text{Might}(\varphi) \wedge \neg\varphi \urcorner$, as in ‘It might be raining and it’s not’. Sentences like this have played an important role in recent work on epistemic modals (see e.g. Groenendijk et al. 1996; Aloni 2000; Yalcin 2007; Dorr & Hawthorne 2013; Ninan 2018; Mandelkern 2019). But the resemblance of practical Moore sentences to these sentences is, as far as I can tell, merely superficial. The infelicity of unembedded sentences with the form of Wittgenstein’s sentence is straightforward to explain, along just the same lines sketched at the outset for accounting for Moore sentences. But, as we have seen, nothing similar extends to practical Moore sentences. Moreover, the discussion of Wittgenstein’s sentences in the recent literature has focused on how to account for the infelicity of these sentences in *embedded* environments; but, as we will see in a moment, practical Moore sentences *embed* in a consistent way, unlike Wittgenstein sentences. A final obstacle for an approach along these lines is that practical Moore sentences sound just as bad when the OPEN conjunct does not contain an epistemic modal, as in $\ulcorner \text{I don’t know whether } \varphi \urcorner$ and so on. But, as is now well known (see citations above), constructions like this do not pattern like epistemic modal constructions in their embedding behavior.

A different, perhaps more promising thought is that the infelicity of practical Moore sentences follows from the *impurity* of modals. Knobe & Szabó (2013) argue that many uses of modals mix “flavors” (see also Phillips & Knobe 2018). One possible explanation of our data, taking its inspiration from this line of thought, would go as follows. In a sentence like ‘You must open the door, even though you might not’, the two modals appear to be interpreted differently: the first is deontic (it concerns what must be done), the second epistemic (it concerns what is known). But if modals are impure, perhaps both modals are interpreted on more of an intermediate basis, which combines both epistemic and deontic flavors. Indeed, perhaps both are naturally interpreted relative to the *same* accessibility relation, in which case—assuming that ‘might’ and ‘must’ are duals—it would be impossible for both to be true. Defenders of this line might then argue that imperatives, performatives, and attitude ascriptions have a core modal meaning which suffices to make them prone to impure modal readings, accounting for all of our data in the same way. That these have a core modal meaning seems quite plausible.¹⁰ It is less obvious to me whether an impure modal approach should be extended to attitudes, imperatives, or performatives, which would go beyond what has to my knowledge been proposed in the literature, but may be plausible.

This approach is even more radical than the one just considered, as it would predict that practical Moore sentences, on the relevant readings, are semantically inconsistent—in other words, that $\lceil \text{ORDER}(\varphi) \rceil$ is inconsistent, not just with $\lceil \neg\varphi \rceil$, but also with $\lceil \text{OPEN}(\neg\varphi) \rceil$. This approach does not seem promising to me. First, it faces the general objections to a semantic approach sketched above: it cannot account for the fact that practical Moore sentences seem knowable; nor can it account for the fact that it can be acceptable to think an OPEN sentence or assert it to someone other than the addressee of one’s order.

This approach also faces two more specific objections. First, as we saw above, variants on practical Moore sentences which contain strong necessity modals or imperatives but which are not used to give orders—i.e., which are used descriptively or permissively, respectively—strike us as felicitous. I do not see how the impure modal account would explain this, since it does not say anything about *ordering* in particular. Without further refinement, it seems to me that this account would predict, for instance, that conjunctions of the form $\lceil \text{Must } p \text{ and I don’t know whether } p \rceil$ would be infelicitous even when the ‘must’ is being used descriptively rather than performatively, since nothing in the account distinguishes the case in which the modal in question is used to give an order from those in which it is used to describe systems of norms.

A second, related obstacle comes from embedding data. Practical Moore sentences do not *embed* like contradictions. Thus for instance the following sound markedly better than the unembedded cases in which they are used to give orders:

- (54) If Liz has to turn in her paper on time but might not do so, we really need to tell her how seriously this is taken here.
- (55) If we don’t know whether Liz will turn in her paper on time but she has to do so, we really need to tell her how seriously this is taken here.

- (56) Suppose that Liz must turn in her paper on time, and that we don't know whether she will do so.
- (57) Suppose I command you to let me go, but I don't know whether you will.
- (58) Mark thinks that Liz has to turn in her paper on time, but he also thinks that she might not do it.

These facts are in line with our generalization that a key source of the infelicity of practical Moore sentences is that they are used to give *orders*; in the embeddings here, the embedded sentence is clearly not being used to give an order, and the result is felicitous, as our generalization predicts. By contrast, this seems to be contrary to the predictions of the impure modal strategy. On that strategy, again, it is not just that $\ulcorner \text{ORDER}(\varphi) \urcorner$ is inconsistent with $\ulcorner \neg\varphi \urcorner$, but rather that $\ulcorner \text{ORDER}(\varphi) \urcorner$ and $\ulcorner \text{OPEN}(\neg\varphi) \urcorner$ are themselves inconsistent. Thus practical Moore sentences would themselves be contradictions, and so are predicted to embed like contradictions. Thus, once more, this approach (without further elaboration) seems to miss a crucial generalization: what is central in generating the infelicity of practical Moore sentences is the use of a modal, imperative, or performative to give an *order*. Other uses—descriptive, permissive, or embedded—do not generate similar infelicity.

In light of these obstacles, an impure modal strategy thus does not seem promising (which is not to say that there are not impure readings of modals in general, just that this does not explain the present data).

In sum, then, I do not think we should look for an explanation of our data by positing a semantic inconsistency between $\ulcorner \text{ORDER}(\varphi) \urcorner$ and $\ulcorner \neg\varphi \urcorner$, or between $\ulcorner \text{ORDER}(\varphi) \urcorner$ and $\ulcorner \text{OPEN}(\neg\varphi) \urcorner$.

4.4 Ninan 2005

A final natural thought is that practical Moore sentences might be explained in the same way as similar sentences discussed by Ninan (2005).¹¹ Ninan points out the infelicity of sentences with the form $\ulcorner \text{ORDER}(\varphi) \wedge \neg\varphi \urcorner$:¹²

- (59) # You must go to confession, but you're not going to.
- (60) # Go to confession! You're not going to go to confession.

Practical Moore sentences are much like Ninan's sentences: schematically, practical Moore sentences are of the form $\ulcorner \text{ORDER}(\varphi) \wedge \text{OPEN}(\neg\varphi) \urcorner$, while Ninan's are of the form $\ulcorner \text{ORDER}(\varphi) \wedge \neg\varphi \urcorner$. The resemblance is more than superficial: any instantiation of Ninan's sentence schema will *entail* corresponding practical Moore sentences. This is because, schematically, φ entails many instances of $\ulcorner \text{OPEN}(\varphi) \urcorner$. For instance, thanks to the factivity of knowledge, φ entails $\ulcorner \text{For all we know, } \varphi \urcorner$. Likewise, on standard approaches, epistemic modals are associated with a reflexive accessibility relation, guaranteeing that φ entails $\ulcorner \text{Might } \varphi \urcorner$. In other words, Ninan's data are strictly stronger variants of my data. This makes it very tempting to look for a common explanation.

I will ultimately propose such an explanation. But Ninan's own explanation won't do the trick. The problem, at a high level, is that Ninan's explanation

crucially exploits the fact that in asserting $\lceil \text{ORDER}(\varphi) \wedge \neg\varphi \rceil$, the speaker commits herself to the falsity of φ . But there is obviously no such commitment in the case of practical Moore sentences. In more detail, Ninan's explanation runs as follows (my emphasis):¹³

In uttering a sentence of the form *must p*, the speaker is undertaking action that she hopes will result in the truth of *p*. . . Thus, if I say, 'You must stop smoking', I am undertaking an action—the action of uttering that sentence—in order to bring it about that you no longer smoke. *Ordinarily, it is rational to undertake an action A to achieve some goal G only if one thinks that there is some chance that one will achieve G if one performs A.* So if I say, 'You must stop smoking' in order to get you to stop smoking, I must think there is some chance that you will stop smoking if I tell you to stop. In that case, I don't believe that you won't stop smoking. But if I don't believe that you won't stop smoking, I'm not in a position to say that you won't stop smoking.

In other words, Ninan takes as background a very natural assumption in action theory:¹⁴ if you are trying to bring it about that φ , you must leave it open that φ will happen. So, if you're trying to make someone do φ , you must leave it open that they will do φ ; i.e., you can't believe that they won't do φ . This straightforwardly accounts for the infelicity of $\lceil \text{ORDER}(\varphi) \wedge \neg\varphi \rceil$. But it is clear that this won't account for our data. The issue is that Ninan's explanation depends on the fact that an assertion of $\lceil \text{ORDER}(\varphi) \wedge \neg\varphi \rceil$ communicates that the speaker believes that φ is false. By contrast, one does *not* have to believe φ is false in order to assert a practical Moore sentence $\lceil \text{ORDER}(\varphi) \wedge \text{OPEN}(\neg\varphi) \rceil$: one need only believe that there is a *possibility* that φ is false. But *that* is not ruled out by the action theoretic principle which Ninan assumes, which entails that you should not order someone to do something you think *will* not happen, but says nothing to rule out ordering someone to do something that you think *might* not happen. Ninan's account thus does not explain our data.

A natural idea at this point is just to strengthen the action theoretic principle which Ninan takes as background so that it covers our cases. Instead of saying that if you're trying to get someone to do φ , you must leave it open that they'll do φ , we could say that if you're trying to get someone to do φ , you have to believe that they *will* do φ . A principle like this would strengthen Ninan's account just enough to cover our data. The problem with this principle is that it is obviously false: cases like the kidnapper case above show that there is nothing wrong with trying to get someone to do φ even when you are not sure they will obey you (and indeed, even when you think there is a rather substantial chance that they will not).

In spite of all this, Ninan's sentences seem closely related to our strictly weaker practical Moore sentences. While Ninan's own account does not extend to our cases, it is natural to pursue a unified explanation of them, as I will do presently.

5. Posturing

Ninan's key principle is one which relates the giving of orders to attitudes: there is something wrong with giving an order to φ if one believes that φ won't obtain.

The natural extension of this to account for our data would say that there is something wrong with giving an order to φ unless one believes that φ *will* obtain. This extension, as we have just seen, is implausible. But there is a more plausible extension which relates ordering to certain attitudes, but in a more subtle way, as follows:

Posturing: When you order someone to φ , you must act towards them as if you believe that they will φ .

I argue that the lesson of practical Moore sentences is that *Posturing* is true. When you order someone to φ , you need not believe that they will φ , but you must act towards them *as if* you believe they will φ . You cannot show weakness by acknowledging that your order might not be obeyed: you must assume an authoritative position vis-à-vis your command.

Let's see first how *Posturing* accounts for our data. When you assert a practical Moore sentence, you are ordering your addressee to φ . You are also making it clear that you are not confident that φ will happen: you are making it manifest that $\lceil \neg\varphi \rceil$ remains a live possibility for you, by conjoining your order with a claim of the form $\lceil \text{OPEN}(\neg\varphi) \rceil$. So you are failing to act towards your addressee as if you believe φ will obtain. If *Posturing* is true, there is something wrong with your speech act.

Posturing also accounts for the contrasts we have seen above between practical Moore sentences and a variety of nearby variants which do not amount to orders. We have seen that close variants on practical Moore sentences which do not amount to orders can be felicitous: this includes variants where a modal is used descriptively or an imperative is used permissively; SHOULD variants; and embedded cases. Since *Posturing* concerns orders in particular, it does not predict there to be anything wrong with these variants—matching observation.

And, insofar as it regulates only the speaker's acts *towards the addressee*, *Posturing* also accounts for another contrast we have seen: the fact that it can be fine to order someone to do something and then assert the corresponding OPEN sentence to someone else—say, to oneself or to some third party. As we saw above, the distinctive infelicity of practical Moore sentences arises from ordering someone to do something and then acting *towards that person* as if they might not—just as *Posturing* predicts.

Another contrast which *Posturing* helps account for, brought to my attention by an anonymous referee for this journal, comes from conditional orders. Conditional orders are orders to do something, conditional on some state of affairs obtaining, as in the following:¹⁵

- (61) a. If you visit Sue, tell her I say hi!
 b. If you visit Sue, you must tell her I say hi!
 c. If you visit Sue, I order you to tell her I say hi!

Since these are not *unconditional* orders, *Posturing* predicts that they can be felicitous in conjunction with *unconditional* expressions of ignorance about whether the

addressee will do the act in question. This seems correct: the sentences in (62) seem fine.

- (62) a. If you visit Sue, tell her I say hi! I don't know whether you will, though, since I'm not sure whether you'll visit her.
 b. If you visit Sue, you must tell her I say hi! Of course, you might not, because you might not visit her at all.

But it is a natural corollary of *Posturing* that conjoining a conditional order with an expression of uncertainty *conditioned on the same proposition as the order* should be infelicitous. This seems correct, as witnessed by the infelicity of the sentences in (63):

- (63) a. # If you visit Sue, tell her I say hi; but if you visit Sue, you might not tell her I say hi!
 b. # If you visit Sue, you must tell her I say hi; but if you visit Sue, you might not tell her I say hi!

Finally, *Posturing* extends immediately to account for Ninan's data. If you assert something of the form $\ulcorner \text{ORDER}(\varphi) \wedge \neg\varphi \urcorner$, you are ordering your interlocutor to do φ while manifestly not acting as if they will. And so, as desired, we have a single explanation of both our data and of Ninan's. (This explanation of Ninan's data is of course consistent with Ninan's; it may well be that both explanations are true, and thus that the infelicity of Ninan's data is overdetermined.)

All this evidence speaks in favor of *Posturing* as an explanation of the infelicity of practical Moore sentences. In the rest of this section, I will say a bit more about the principle at a higher level of abstraction. *Posturing* draws on two strategies which we explored and rejected in §4. The first is the idea that what is wrong with our infelicitous conjunctions is broadly rhetorical: why talk about your uncertainty about whether someone will do something when you're telling them to do it? As we saw above, this idea is too broad: it does not predict the contrast between our order-giving practical Moore sentences and nearby variants which are not used to give orders. The second is the extension of Ninan's action theoretic principle, which, as we saw, would cover our cases but is simply not plausible. *Posturing* makes the rhetorical strategy more discriminating, and grounds it in a more plausible action theoretic principle: what is bad is not giving advice without acting towards your addressee as though you are confident the advice will be followed; what is bad is rather specifically *ordering* your addressee to do φ without acting towards them as though you are confident they will. And the badness has to do with something general about ordering: to acknowledge that your interlocutor might not obey you is to acknowledge that you are in a position of some weakness with respect to your act of ordering. And, it seems, we are disinclined to acknowledge such weakness, even when it is manifestly present—indeed, even when it is mutually acknowledged, as in the kidnapping case. We are obligated, as it were, to put a good face on things, and act as though we are in full control of our interlocutor, even when it is clear that we are not.

Posturing is fascinating as a thesis about human psychology, and there is much work to do in situating it within a broader theory of command and authority. I will not try to do much of that work here; my main goal here has been rather to explain the puzzle that practical Moore sentences pose, and provide a preliminary solution to that puzzle. I suspect that a full account of *Posturing* will connect it to the way that we think about authority in general, and I will say just a bit about how that connection might look. Consider the following possible explanation of *Posturing*. In the literature on authority, it has been suggested that to fully exercise authority over someone is to give that person practical reasons for action. Here is a characteristic formulation: ‘The commander characteristically intends his hearer to take the commander’s will instead of his own as a guide to action’ (Hart, 1982).¹⁶ And it has been proposed by action theorists that when you intend to do something you must believe you will do it.¹⁷ Putting these thoughts together, it follows that for *A* to fully exercise authority over *B* in ordering *B* to φ , *A* is, essentially, intending that φ happen *via B’s* agency, and thus *A* must believe that φ will happen. Now we have seen that it not plausible that, when you give a command, you must believe it will be obeyed. But a more plausible thought is in the neighborhood: we could say that one can give a command without having *full authority* over one’s addressee in the sense just sketched; but that, in giving a command, one must *act as if* one has full authority over one’s addressee. Putting all this together, we arrive at *Posturing*: in giving a command to φ , you must act towards your addressee as if you have full authority over them, which in turn requires acting as though you believe they will actualize your intention that they φ .

Every step of this story is controversial, and I will not commit myself to any of it here. I am not sure how much of it I myself believe; in particular, I am skeptical of the principle that if you intend to do something, you must believe you will do it. My goal with this sketch is simply to give a sense of the kind of explanation we might find of *Posturing* by way of a more careful exploration of broad principles in action theory and the theory of authority. Such considerations will help us understand *Posturing*; and, of course, *Posturing* may well help us decide between rival theories of authority and orders, since, if the main argument in this paper is correct, our theory of these will have to account for *Posturing*.¹⁸

This discussion assumes that *Posturing* has a fairly deep psychological foundation. We could alternately imagine it resting on a local, conventional foundation, having to do with the particular way we happen to think about orders. Then violations of *Posturing* would be like an officiant saying ‘I now pronounce you man and wife, although I have no legal authority to do so’—a speech act which is incoherent only because of local and contingent facts about marriage rites.¹⁹ But it strikes me as likely that *Posturing* has a deeper foundation than that. This is because *Posturing* concerns ordering, which—while it surely has diverse attendant conventions in different societies—strikes me as more likely than not to be in itself a cross-linguistically and cross-culturally robust category of speech act (very much unlike, for instance, the speech act of marrying). This raises the question, in turn, of whether *Posturing* obtains across different languages and cultures. I will not try to answer this question. If the answer turns out to be affirmative, however, that

would show that certain apparently *rhetorical* rules which might seem elective are, in fact, hardwired. The advice given in *Posturing* might seem like reasonable advice to give someone about how to appear more authoritative; we might even imagine it being backed up by psychological research which shows that people will be more likely to follow your orders if you act towards them as if you are confident that they will. But it strikes me as surprising that something like *Posturing* should be hardwired as a universal requirement on orders. In other words, what we would then have would be something which *prima facie* looks like nothing more than a good rhetorical rule of thumb, but which turned out to be encoded in something like linguistic knowledge. Now, whether or not this actually *is* what we have here is a question which will turn, again, on whether practical Moore sentences strike speakers as infelicitous across different languages and cultures. If it does, it would show that apparently rhetorical rules like *Posturing* can be, to a surprising degree, built into our understanding of language and the nature of certain fundamental speech acts.

6. Objections and Replies

Before concluding, let me consider a few possible objections to the account I have given.

6.1 Third-personal variants

The first objection comes from third-personal variants on our practical Moore sentences, like (64):

(64) Liz must turn in her paper on time, but she might not.

What should we say about sentences like (64)? The first question to ask is whether such sentences are felicitous. It seems to depend on the context. In contexts where it is clear that the first conjunct in (64) is not a command, but simply a description of the requirements that bind Liz, it seems to be felicitous. Suppose that two TAs are speculating on whether Liz will pass the class. One says:

(65) Well, I'm not sure she will. She has to turn her paper in on time, but she might not get it done. Then she would fail.

Here it is clear that the TA is not giving a command to Liz of any kind; she is simply describing the constraints that bind Liz. And in this case, the conjunction in (64) sounds perfectly felicitous, in line with our observation at the outset that the infelicity of practical Moore sentences arises from the fact that they are used to give *orders*.

But there are other contexts where (64) sounds infelicitous, and these contexts constitute a *prima facie* challenge for our approach. Suppose Sue says (64) to her TA, who has asked Sue whether Liz, a chronic procrastinator, can have an extension on her final paper. In that context, (64) sounds very odd. This is a *prima facie* challenge for our account. One might think that all orders are second-personal,

and thus that (64) cannot be an order. In that case, *Posturing* would not have anything to say to explain why (64) is infelicitous in the relevant case.

I think what we should say about cases like this is that there are third-personal orders, and that, in the relevant context, (64) amounts to exactly that. Although Liz is not present when (64) is asserted, it is still reasonable to interpret (64) as an order which requires Liz to turn her paper in on time.²⁰ If we countenance third-personal orders, as I believe we should, then our account rightly predicts (64) will be infelicitous in this case; that this is the right account is suggested by the contrast with the first context, where (64) is clearly not an order, and sounds felicitous.

6.2 *Embeddings*

The account I have given relies essentially on the *speech act* properties of practical Moore sentences. I have adduced a variety of evidence that suggests that this is the right level of explanation. However, if the infelicity of practical Moore sentences persists in embedding environments where they are not (even conditional) orders, that would be a challenge for this kind of approach; conversely, if the infelicity goes away, that would provide support for my account. To see this point, it is useful to compare the parallel situation for Moore sentences: if the standard account of Moore sentences is correct, then the infelicity of Moore sentences will not percolate up to sentences which embed them in entailment-canceling environments. This turns out to be exactly how things work: thus e.g. ‘If it’s raining and I don’t believe it’s raining, then I’m being an idiot’ is perfectly coherent.

Crucially, as we saw in §4.3 above, this is how things work for practical Moore sentences as well: their infelicity does not percolate up to sentences which embed them in entailment-canceling environments. This provides further support for an analysis of the phenomenon at the level of speech acts.

6.3 *Conditionals*

The third objection, suggested to me by Daniel Rothschild, comes from the interplay of orders and conditionals (Silk (2018, Fn. 36) discusses similar data). Sentences like the following are felicitous:

- (66) You must turn your paper in on time. If you don’t turn it in on time, you’ll fail the class.

This is surprising from the point of view of *Posturing* given standard assumptions about indicative conditionals. Indicative conditionals are often taken to entail or presuppose that the speaker leaves open the possibility that the antecedent is true (see Stalnaker 1975; von Stechow 1998 for discussion). On this line, ‘If φ , ψ ’ entails or presupposes ‘ $\text{OPEN}(\varphi)$ ’, which means that (66) would license an inference to the practical Moore sentences ‘You must turn your paper in on time, but you might not’. But then we would expect (66) to be infelicitous, given *posturing*.

The underlying theoretical issues here are quite complex, but there is a simple way to defuse this objection: by noting the felicity of conditionals like (67).

(67) Liz will turn her paper in on time. [And] if she doesn't, I'll be furious.

If 'If she doesn't [turn her paper in on time], I'll be furious' licensed the inference to 'Liz might not turn her paper in on time', then (67) would license the inference to:

(68) # Liz will turn her paper in on time. [And] she might not turn her paper in on time.

But it clearly does not: (68) sounds awful, while (67) does not.

This suggests that the usual OPEN inferences of indicative conditionals are somehow blocked in conditionals like the one in (67). I do not know of discussion of this phenomenon, and I do not know how to make sense of it. But this observation suffices to defuse the current objection: it looks like conditionals like those in (66) and (67) do not, for some reason, in general license the inference to the corresponding OPEN sentence. There is interesting variation when we have an explicit future orientation in the antecedent of the conditional: thus for instance (69) is substantially worse than (67).

(69) # Liz will turn her paper in on time. If she isn't going to, I'll be furious.

Conditionals like this also clash with the corresponding order:

(70) # You must turn your paper in on time. If you're not going to, you'll fail the class.

Again, I do not know what to make of this pattern in general. But the fact that the infelicity of \ulcorner ORDER(φ). If $\neg\varphi, \psi \urcorner$ patterns with the infelicity of $\ulcorner\varphi$. If $\neg\varphi, \psi \urcorner$ is exactly in line with the predictions of *Posturing*.²¹

7. Conclusion

I have described a practical analogue of Moore sentences: speech acts in which one gives an order, while acknowledging that one is less than certain that it will be obeyed. The infelicity of practical Moore sentences is surprising. I have offered an account of their infelicity on the basis of the *Posturing* principle, which says that, when you order someone to do something, you must adopt a posture of confidence towards them: you must act towards them as though you are confident that they will comply. This explains the infelicity of practical Moore sentences, while accounting for the fact that close variants which do not amount to orders do not exhibit the same infelicity. If this account is correct, it is a striking fact which may teach us about the norms governing authority and action, as well as about the ways that these norms are cognitively encoded.

My discussion here has been necessarily preliminary. Further exploration of the data; of alternate possible solutions; and of the foundations of my proposed solution are in order. We should, for instance, explore whether we find evidence

for *Posturing* in other linguistic contexts.²² We should also look for evidence in non-linguistic domains. *Posturing* is not a specifically linguistic norm, and so, if it is correct, we should find evidence for it beyond the linguistic domain: in particular, *Posturing* predicts that giving an order and then acting non-verbally in a way that makes it clear to your addressee that you leave it open that you will not be obeyed will be felt to be strange. If this prediction were not borne out, then we would have to conclude that in place of *Posturing* we need a more specifically *linguistic* norm. Settling this question will, in turn, help determine what kind of foundation we should seek for the relevant norm—whether I have been right to tentatively suggest that the foundation lies in the nature of authority and ordering broadly speaking, or whether instead we should look for a more narrowly linguistic explanation.

We should, finally, explore the ramifications of *Posturing* for broader questions about the norms governing speech acts. To give a sense of the kind of light *Posturing* might shed here, consider the following puzzle. It is commonly accepted that one should know, or at least believe, what one asserts. A norm of this kind can be motivated on the basis of Moore sentences like those we saw at the outset. But there are variants of Moore sentences which sound equally bad, but are not explained by such norms, like (73):

(73) # It's raining, but I'm not absolutely certain that it's raining.

The infelicity of (73) is not explained straightforwardly by a knowledge or belief norm on assertion, since a speaker can believe, and indeed know, (73) (provided that knowledge does not require absolute certainty). (73) *would* be explained by a norm which said one must *know with absolutely certainty* whatever one asserts. But it is intuitively clear there is no such norm: one can assert in the absence of absolute certainty. *Posturing* may offer a helpful lens through which to understand this situation. Perhaps what is required here is not that you know with certainty what you assert, but rather that you *act as if* you do. This would be far more plausible than a norm which says you must know with certainty what you assert, and would account for the infelicity of (73). There thus may be a whole bevy of norms which, like *Posturing*, require not that one have a certain cognitive relation towards a content but rather that one *act as if* one has that cognitive relation.²³

Notes

¹ See Hintikka 1962 for this line of reasoning. Wittgenstein's variant with 'might' in (4) introduces some interesting complications, since sentences like (4) embed differently from Moore sentences; see §4.3 for discussion.

² I assume that 'but' and 'and' both have the core semantics of conjunction, and schematize both using '∧'.

³ Silk (2015)'s example is 'I'm not sure if Alice will come clean, but she must'. Silk's discussion is brief, but, as will become clear, I am in agreement with the basic idea: 'Uttering "Must φ " not only conveys that $\ulcorner \neg\varphi \urcorner$ -possibilities are unacceptable; it suggests that $\ulcorner \neg\varphi \urcorner$ -possibilities aren't even on the table for consideration. Of course obligations can go unfulfilled. What is interesting is that speakers appear to assume otherwise, at least for the purposes of conversation, when expressing obligations with 'must'.' Ninan (2005) discusses similar but importantly different sentences; see §4.4 for discussion.

Frank (1996); Kaufmann (2012); Condoravdi & Lauer (2012); Harris (2017) discuss ‘deontic Moore paradoxes’, which are conjunctions like ‘Call him at home! I don’t want you to’; these are different from our practical Moore sentences, and, I think, pose different problems.

⁴ That these are used to give orders is obvious for the imperative and performative variants. That this is so for the strong deontic necessity modals (‘must’, ‘have’) is less obvious, but something that Ninan (2005) argues for, citing Leech 1971; Kamp 1973; Coates 1983; Palmer 1986; I will follow Ninan here, making this natural assumption without further argument. I do not claim that imperatives or strong deontic necessity modals are *only* used to give orders, as I discuss further presently.

⁵ For simplicity, I will assume we can model actions as propositions. I will be sloppy about use and mention, and I will move somewhat freely between thinking of schemas like ‘ORDER(φ)’ as ranging over uses of sentences versus the sentences used.

⁶ See discussion in Ninan 2005; von Stechow & Iatridou 2008 and citations therein.

⁷ There are a variety of ways of accounting for this; what’s crucial for my purposes is simply that neither imperative has the force of an order to do its prejacent.

⁸ An anonymous referee for this journal points to another variant on practical Moore sentences that seem to be potentially felicitous: cases where the speaker gives an order while simultaneously distancing herself from the order she is giving, as in the following.

(40) You must close the door. I don’t know if you will, but I must order you to do so anyway.

(41) Close the door! I know you may not comply with this command, but it is my duty to issue it nevertheless.

These seem to be somewhat more felicitous than the cases considered above. One way of making sense of this is that, in these cases, the speaker is distinguishing two roles in which she stands: one role, *qua* issuer of a command, is as a representative of some relevant institution/system of norms (whatever it is that puts her in a position to be obligated to issue the relevant command); the other role, *qua* expresser of uncertainty, is as an individual agent independent of the relevant institution. Insofar as she is switching roles between the giving of the command and the expression of uncertainty, these seem to be felicitous. This suggests that a key source of infelicity in our core cases is the *simultaneous* giving of an order while expressing uncertainty.

⁹ Thanks to an anonymous referee for this journal for helpful discussion concerning this idea.

¹⁰ See Kaufmann 2012 for this approach to imperatives, and of course the classic approach to attitudes due to Hintikka 1962 treats them as essentially modal operators.

¹¹ Ninan is also concerned with the distribution of deontic versus epistemic readings over different tenses, and aims to give a unified account of both issues. I will not address that second issue.

¹² See related discussion in Palmer 1990; Kaufmann 2012.

¹³ Ninan first discusses and then rejects a different solution. On that line, $\lceil \text{Must } \varphi \rceil$ is only felicitous when φ and $\lceil \neg\varphi \rceil$ are both compatible with the speaker’s knowledge. If we could motivate such a constraint in general, this would account for Ninan’s data, but it obviously won’t help with our data.

¹⁴ Citing Bratman 1987.

¹⁵ See Schwager 2006; Kaufmann & Schwager 2009; Harris 2017 for recent discussion.

¹⁶ See also e.g. Raz 1975a, 1975b, 1988; Roth 2004; Hanser 2015 for related discussion.

¹⁷ See e.g. Grice 1972; Davidson 1980; Harman 1986; Bratman 1987; Velleman 1989; Portmore 2019 for arguments in favor and against.

¹⁸ Thanks to Hasan Dindjer and Brendan de Kenessey for very helpful discussion on these points.

¹⁹ Thanks to an anonymous referee for this journal for this example and for helpful discussion of these issues.

²⁰ Ninan (2005) advocates a similar reply in response to a parallel objection to his account. The plausibility that there can be third-personal orders seems to me to be increased by the existence of languages which have grammatical imperatives in the third person.

²¹ A similar sort of worry, raised by Hasan Dindjer (p.c.), comes from orders combined with questions. ‘You must close the door! Will you?’ If questions are appropriate only when the speaker does not know the answer, these pose a problem for *Posturing*. I am inclined to think, however, that in this

case, the question ‘Will you?’ is not in the first instance expressing ignorance but rather is seeking verbal consent to the order.

²²As a referee for this journal helpfully points out, *Posturing* predicts in particular that speech acts which give an order and express the corresponding OPEN sentence will generally be infelicitous, even if the order is given using language which is not canonically used to give orders. For instance, following that referee’s suggestion, consider a context where it is accepted that any claim made with a weak necessity will be treated as an order. In such a context, *Posturing* predicts that a sentence like (71) will be infelicitous:

(71) You should close the door, but you might not.

I find it hard to fix clearly on intuitions about sentences like (71) in the relevant context. One way to argue that the prediction of *Posturing* here is correct is that if we add some material making explicit that (71) is an order, it becomes more clearly infelicitous:

(72) # You should close the door (and don’t forget: that’s an order!), but you might not.

²³Thanks to Daniel Rothschild for suggesting this possibility.

References

- Aloni, Maria. (2000). Conceptual covers in dynamic semantics. In Lawrence Cavedon, Patrick Blackburn, Nick Braisby & Atsushi Shimojima (eds.), *Logic, language and computation*, vol. III, 23–48. Stanford, CA: CSLI: Language and Computation.
- Coates, J. (1983). *The semantics of modal auxiliaries*. London and Canberra: Croom Helm.
- Condoravdi, Cleo & Sven Lauer. (2012). Imperatives: Meaning and illocutionary force. In Christopher Piñón (ed.), *Empirical issues in syntax and semantics*, vol. 9, 37–58.
- Davidson, Donald. (1980). Intending. In *Essays on actions and events*, Oxford: Oxford University Press.
- Dorr, Cian & John Hawthorne. (2013). Embedding epistemic modals. *Mind* 122(488). 867–913. 10.1093/mind/fzt091.
- von Fintel, Kai. (1998). The presupposition of subjunctive conditionals. In Orin Percus & Uli Sauerland (eds.), *The interpretive tract*, vol. 25, 29–44. Cambridge, Massachusetts: MIT Working Papers in Linguistics. <https://philpapers.org/rec/VONTPO-2>.
- von Fintel, Kai & Sabine Iatridou. (2008). How to say ‘ought’ in foreign: The composition of weak necessity modals. In J. Guéron & J. Lecarme (eds.), *Time and modality*, 115–141. Springer.
- von Fintel, Kai & Sabine Iatridou. (2017). A modest proposal for the meaning of imperatives. In Ana Arregui, María Luisa Rivero & Andrés Salanova (eds.), *Modality across syntactic categories*, Oxford University Press. 10.1093/acprof:oso/9780198718208.003.0013.
- Frank, Annette. (1996). *Context dependence in modal constructions*: Universität Stuttgart dissertation.
- Grice, H. Paul. (1972). Intention and uncertainty. In *Proceedings of the British Academy*, vol. 57, 263–279.
- Groenendijk, Jeroen, Martin Stokhof & Frank Veltman. (1996). Coreference and modality. In Shalom Lappin (ed.), *Handbook of contemporary semantic theory*, 179–216. Oxford: Blackwell.
- Hanser, Matthew. (2015). Doing another’s bidding. In George Pavlokos & Veronica Rodriguez-Blanco (eds.), *Reasons and intentions in law and practical agency*, 95–120. Cambridge University Press.
- Harman, Gilbert. (1986). *Change in view*. Cambridge, MA: MIT Press.
- Harris, Daniel W. (2017). Imperatives and intention-based semantics. <https://semanticsarchive.net/Archive/mFiNjhmN/DanielWHarris-ImperativesAndIntentionBasedSemantics.pdf>. Manuscript, Hunter College.
- Hart, H.L.A., (1982). *Essays on Bentham: Jurisprudence and political philosophy*. Oxford University Press.
- Hintikka, Jaakko, (1962). *Knowledge and belief: An introduction to the logic of two notions*. Ithaca, NY: Cornell University Press.
- Kamp, Hans. (1973). Free choice permission. In *Proceedings of the Aristotelian Society*, vol. 74, 57–74.

- Kaufmann, Magdalena. (2012). Interpreting imperatives. In *Studies in linguistics and philosophy*, vol. 88, New York: Springer.
- Kaufmann, Stefan & Magdalena Schwager. (2009). A unified analysis of conditional imperatives. In Ed Cormany, Satoshi Ito & David Lutz (eds.), *Semantics and Linguistic Theory (SALT)*, vol. 19, 239–256.
- Knobe, Joshua & Zoltán Gendler Szabó. (2013). Modals with a Taste of the Deontic. *Semantics and Pragmatics* 6(1). 1–42. 10.3765/sp.6.1.
- Leech, G., (1971). *Meaning and the English verb*. Harlow, UK: Longman.
- Lewis, David. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic* 8(1). 339–359.
- Mandelkern, Matthew. (2019). Bounded modality. *The Philosophical Review* 128(1). 1–61. 10.1215/00318108-7213001.
- Michael, Bratman. (1987). *Intention, plans, and practical reason*. Cambridge, MA: Harvard University Press.
- Moore, G.E. (1942). A reply to my critics. In P.A. Schilpp (ed.), *The philosophy of G.E. Moore*, Evanston, IL: Northwestern University.
- Ninan, Dilip. (2005). Two puzzles about deontic necessity. In J. Gajewski, V. Hacquard, B. Nickel & S. Yalcin (eds.), *New work on modality*, vol. 51, 149–178. MIT Working Papers in Linguistics.
- Ninan, Dilip. (2018). Quantification and epistemic modality. *The Philosophical Review* 127(2). 433–485. 10.1215/00318108-6973010.
- Palmer, F., (1986). *Mood and modality*. Cambridge, UK: Cambridge University Press.
- Palmer, F. R., (1990). *Modality and the English modals*. New York: Longman.
- Phillips, Jonathan & Joshua Knobe. (2018). The psychological representation of modality. *Mind and Language* 33(1). 65–94. 10.1111/mila.12165.
- Portmore, Douglas W., (2019). *Opting for the best*. Oxford: Oxford University Press.
- Raz, Joseph. 1975a. *Practical reason and norms*. Oxford University Press.
- Raz, Joseph. 1975b. Reasons for action, decisions and norms. *Mind* 84(336). 481–499.
- Raz, Joseph. (1988). The justification of authority. In *The morality of freedom*, 38–69. Oxford University Press.
- Roth, Abraham Sesshu. (2004). Shared agency and contralateral commitments. *The Philosophical Review* 113(3). 359–410. 10.1215/00318108-113-3-359.
- Schwager, Magdalena. (2006). Conditionalized imperatives. In C. Tancredi, M. Kanazawa, I. Imani & K. Kusumoto (eds.), *Semantics and Linguistic Theory (SALT)*, vol. 16, 241–258.
- Silk, Alex. (2015). What normative terms mean and why it matters for ethical theory. In M. Timmons (ed.), *Oxford studies in normative ethics*, vol. 5, 296–325. Oxford University Press. 10.1093/acprof:oso/9780198744665.003.0014.
- Silk, Alex. (2018). Weak and strong necessity modals. In B. Dunaway & D. Plunkett (eds.), *Meaning, decision, and norms: Themes from the work of Allan Gibbard*, Maize Books: Michigan Publishing, University of Michigan.
- Stalnaker, Robert. (1975). Indicative conditionals. *Philosophia* 5(3). 269–86. <https://philpapers.org/rec/STAIC>.
- Velleman, David, (1989). *Practical reflection*. Princeton University Press.
- Wittgenstein, Ludwig. 2001/1953. *Philosophical investigations*. Wiley-Blackwell 3rd edn.
- Yalcin, Seth. (2007). Epistemic modals. *Mind* 116(464). 983–1026. 10.1093/mind/fzm983.