

11

Permission and (So-Called Epistemic) Possibility

Stephen Yablo

It is just possible that I can explain the structure of this paper by comparing it to a Soviet-era joke, based apparently on a remark of John Kenneth Galbraith's.

Worker to party official: Tell me, what is the difference between Capitalism and Communism?

Party official: Good question, comrade. It's this. Under Capitalism, you see, man exploits man. Under Communism, it's the other way around.

The joke operates on several levels. "The other way around" presumably means that under communism, *man* exploits *man*; and that sounds exactly the same as what was said about capitalism. Second, though, it's not as though we can't hear a difference between the first "man exploits man" and the second; one has the sense it's different men in the two cases. Third, though, that difference overlays a deeper similarity; one group exploiting another makes for a lot of the same unpleasantness as the other group exploiting the first.

Now consider the sentence, or sentence-fragment, *that may be*. It too can be read in two ways. On the one reading, it indicates that a certain state of affairs, φ let's call it, is not ruled out. On the other reading, it *also* indicates that φ is not ruled out.

This is a distinction? Just as in the exploitation joke, it's hard to see what difference there could be between φ 's not being ruled out and, well, φ 's not being ruled out. But, and this is the second point of analogy, it's not as though we can't hear a difference between different utterances of "Such and such is not ruled out". Compare "Sabotage is not ruled out," said by an FAA investigator after

I am grateful to Seth Yalcin, Eric Swanson, Bob Stalnaker, Thony Gillies, Kai von Fintel, Caspar Hare, Carrie Jenkins, Kit Fine, George Bealer, Andy Egan, Tim Williamson, Josh Dever, Nicholas Asher, David Beaver, Cleo Condoravdi, Alejandro Pérez Carballo, Paolo Santorio, and participants (beyond those already mentioned) in the Arché Conference on Modality (St Andrews, 2006), especially my excellent commentator David Eford, and the Austin Workshop on Epistemic Modals (University of Texas at Austin, 2006).

the plane crashes, and “Sabotage is not ruled out” said by a rebel leader six months before the crash. The investigator is saying that sabotage is not ruled out *descriptively*, as it would be if she’d asserted, “There was no sabotage”. The rebel leader is saying that it is not ruled out *prescriptively*, as it would be if she’d commanded her underlings not to engage in sabotage. Yet, and this is the third point of analogy, just as the two things it can mean for man to dominate man have at a deeper level a lot in common, the two things “That may be” can mean have at a deeper level a lot in common. The descriptive reading, “That may be *so*,” has a lot in common with the deontic reading, “That may be *done*.” That anyway is what I’ll be arguing in this paper.

What does it matter, though, if the one “may” has properties in common with the other? It matters because descriptive (often called *epistemic*) “may” is *extremely confusing*; and the questions we are driven to as we attempt to understand it are questions that, as it happens, have been much discussed in connection with deontic “may”.

The standard semantics for “it may be (or might be, or it is possible) that φ ”¹ has it expressing something in the vicinity of the speaker’s failing to know that $\sim\varphi$. Thus Moore:

It’s possible that I’m not sitting down now . . . means ‘It is not certain that I am’ or ‘I don’t know that I am’²

More sophisticated versions allow the knower(s) and/or the information against which φ is tested to vary:³

It is possible_A that p is true if and only if what A knows does not, in a manner that is obvious to A , entail not- p .⁴

“It is possible that p ” is true if and only if (1) no member of the relevant community knows that p is false, and (2) there is no relevant way by which members of the relevant community can come to know that p is false.⁵

There is undoubtedly something right about the standard semantics. But there are things wrong with it too.

One problem is that it gets the subject matter wrong. When I say, “Bob might be in his office,” I am talking about Bob and his office, not myself or the extent of my information.⁶ Suppose the building is on fire and everyone has been told

¹ I will generally use “might” rather than “may” for so-called epistemic possibility.

² Moore (1962), p. 184.

³ DeRose (1991); Hacking (1967); Teller (1972); von Fintel and Gillies (2008).

⁴ Stanley (2005), p. 128. ⁵ DeRose (1991), pp. 593–4.

⁶ It might seem this worry could be sidestepped by putting speaker’s knowledge into the mechanism by which the content is generated, rather than the content itself (as in Kratzer 1981). Speaker’s knowledge would play the same sort of role in the evaluation of “might”-claims as speaker’s attention plays in the evaluation of “you”-claims. Andy Egan has convinced me that

to leave. I am afraid that Bob might still be in his office. I am not afraid that I don't know he's elsewhere.

Two, the proposed truth-conditions, at least in their naïve Moorean form, are too weak. The mere fact that I don't myself know that $\sim\varphi$ doesn't make it true in my mouth that φ might be so. Suppose you question my claim on the basis that Bob was just seen stepping onto a plane. It would be no reply at all to say that my information really was as limited as I suggested; I really and truly didn't know that Bob was not in his office. Evidently the information that needs to comport with φ for a might-claim to be correct can extend beyond what the speaker personally knows at the time of utterance. The proposal should really be that "it might be that φ " is true iff φ is not ruled out by any *pertinent* facts—where the test of pertinence is presumably that the speaker is prepared to acknowledge that she was mistaken if these facts really do/did obtain.⁷

But, and this is the third problem, the truth-conditions are now so strong that speakers generally have no idea whether they are satisfied—so strong that speakers should refrain from asserting that φ might be so. The principle here is that I should not assert that χ , if (a) I am aware of a ψ such that χ is false if ψ is true, and (b) I consider ψ entirely likely to be true. When χ is "it might be that φ ", I am virtually always aware of a ψ like that, viz. "somewhere out there, there is evidence that rules φ out". ψ meets condition (a) because I freely accept that my might-claim is mistaken if φ is ruled out by the evidence, including evidence I don't myself possess. I freely accept, for instance, that if Bob was seen getting on a plane at 11:55, then it is not true that Bob might now (at noon) be in his office. ψ satisfies condition (b) because I do not, when I say that it might be that φ , take myself to know that my evidence is relevantly complete; *obviously* I might be missing something which makes φ highly unlikely.⁸ It is clear to me, when I say that Bob might be in his office, that evidence not in my possession might well show he is somewhere else. But now, if I think it entirely possible that there is evidence that exposes my statement as false, how in good conscience can I make the statement? Who would dare make a might-claim, if the claim was entirely likely to be wrong?

The fourth problem with the standard semantics is that it is too *epistemic*. I have a thing about the sanctity of the ballot box, imagine, so when you ask me whether I am going to vote for Kucinich, I say, "I might or I might not," despite knowing perfectly well what I've decided, and not trying to hide the fact that I

there is still a *prima facie* problem, since the Kratzer-proposition is at root a consistency claim, and consistency-claims have the wrong evidential and modal properties.

⁷ For more on these issues, see Egan (2007); Egan, Hawthorne, and Weatherson (2005); MacFarlane (2003, forthcoming); von Fintel and Gillies (2008).

⁸ If that kind of knowledge were required, I would not say "it might be that φ and it might be that $\sim\varphi$ " unless I thought that φ was *objectively undecidable*, in the sense that all the evidence in the world left it an open question whether φ .

know. (I am lying if I say I don't know whether I'll vote for Kucinich, but it is not a lie to say I might vote for him or I might not.) Or suppose that I run into a creditor who demands that I give him a check by the end of the day. I know perfectly well that I am going to do what he asks—I have the check in my pocket—but still I say, “I might or I might not; it might have to wait until tomorrow,” for the loan is not strictly due until Friday. I say it not because I don't know I'll give him the check today, but because I reserve the right not to; it's not a limit in my information I'm indicating, but rather a limit in what I'm prepared to commit to. One final example. Imagine I am pitching a story line to a Hollywood mogul. “Now comes the good part,” I tell him. “The Raskolnikov character brutally murders the pawnbroker.” “Not a chance, not if we want PG-13,” comes the reply. “OK,” I say in a concessive spirit, “so he might just rough her up a bit.” The “might” here is not to indicate the limits of my knowledge; there is nothing to know in this case, since the movie almost certainly won't get made. It's to indicate the limits of my proposal.

A fifth problem is more logical in nature. Suppose φ is consistent with all pertinent information. Then so is everything φ entails. One would expect, then, that if φ entailed ψ , “it might be that φ ” would entail “it might be that ψ ”. One would expect, for instance, that “Bob might be in his office” would entail “Bob might be in his office or in an opium den”. And yet “Bob might be in his office or in an opium den” makes a stronger claim, to the effect that Bob might be in the one place *and* in addition he might be in the other. There is of course a similar puzzle about permission: how is it that “you can go or stay” entails (or seems to) that you can do whichever you want, that is, it is open to you to go *and* it is open to you to stay?⁹

The sixth problem I learned from Seth Yalcin. An advantage sometimes claimed for the standard semantics is that it explains the paradoxicality of “ φ and it might be that $\sim\varphi$ ”. To say that is to say in effect, “ φ but I don't know that φ ”. And that's just Moore's paradox. The usual line on Moore's paradox is that the Moore-sentence is *unassertable* even though there's no reason it can't be true. The proof it's a problem of assertability rather than truth is that there's nothing to stop me from *supposing*, say in the antecedent of a conditional, that φ and I don't know it; “if a piano is unbeknownst to me falling on my head, I am in for a big shock.” One would expect, then, that “ φ & it might be that $\sim\varphi$ ” would be supposable too. And it isn't. “If a piano is falling on my head but it might not be falling on my head, . . .” makes no sense. The sixth problem is that “ φ & it might be that $\sim\varphi$ ” is not coherently supposable, and the standard semantics offers no explanation of this.

So the traditional “static” semantics for epistemic modals has its problems. This has led some to propose a *dynamic* semantics; the meaning of “might φ ” is given

⁹ Kamp (1973); Zimmermann (2000).

not by its truth-conditions but its effect on context or shared information. The most popular version of this is Frank Veltman's *default semantics*.¹⁰ Veltman says that "might φ " uttered in information state S returns that same information state S if S is consistent with φ , and returns the null information state if S is not consistent with φ . Both parts of this seem *prima facie* at odds with the way "might" is used.

Consider first the idea that "might φ " returns the null state when S is inconsistent with φ . Suppose it is understood all around that John, Paul, George, and Ringo will be at the party. Then Yoko runs in with the news that Ringo might not be able to make it. Ringo's not making it is inconsistent with John, Paul, George, and Ringo being there. All our information is demolished, then. Intuitively, though, the information that John, Paul, and George will be at the party remains when we learn that Ringo might not attend.¹¹

The idea that "might φ " returns S if φ is consistent with S seems questionable, too. It's consistent with John, Paul, George, and Ringo being at the party that Ringo or Elton John stays away; for it might be Elton John that stays away. Nevertheless, if Yoko runs in with the news that Ringo or Elton John might not be attending, we will hardly keep on assuming that all four Beatles will be there. Rather our shared information is weakened to: John, Paul, and George will be at the party.

It seems from these examples that "might φ ", uttered in information state S, has or can have the effect of cutting S down to a weaker information state S'; and it can have this effect both when φ is consistent with S and when φ is inconsistent with S. If we model information states with sets of worlds, then the effect of "might φ " is to add on additional worlds. The question, of course, is *which* additional worlds. The reason for looking at *deontic* modals is that the analogous question about them was raised years ago by David Lewis, in a paper called "A Problem about Permission".

Lewis starts by describing a simple language game. The players are Master, Slave, and Kibitzer, though we'll be ignoring Kibitzer (he's used to it). Master issues commands and permissions to Slave, thereby shrinking and expanding what Lewis calls the *sphere of permissibility*, the set of worlds where Slave behaves as he's supposed to. Behaving as he's supposed to is Slave's only purpose in this game, and given how we defined the sphere of permissibility, that comes to behaving so that the actual world lies within that sphere. Slave can't stay within the sphere, though, unless he knows where the sphere is. Let's try to help him with this: how does the sphere evolve?

¹⁰ Gillies (2004); Veltman (1996).

¹¹ The counter-response is that we would take steps to avoid this disaster by scaling back to an information state consistent with Ringo's non-attendance. I agree that this is what we would do; the question is whether scaling back should be understood as a repair strategy used when disaster threatens, or as part of "might"'s basic semantic functioning. See Fuhrmann (1999). Thanks here to Thony Gillies.

When the game begins, all worlds are permissible. Now Master begins issuing commands and permissions. Our job is to figure out the function that takes a given sequence of commands (written $!\varphi$) and permissions (written $i\varphi$) to the set of worlds permissible after all those commands and permissions have been given. That fortunately boils down to two simpler-seeming sub-tasks: first, figure out the effect of a *command* on the sphere of permissibility, second, figure out the effect of a *permission* on the sphere of permissibility.

You might think the second sub-task would be the easier one: a sphere of *permissibility* would seem to be more directly responsive to *permissions* than commands. But it's actually the first sub-task that's easier. Suppose the going sphere of permissibility is S and Master says "Mop that floor!" Then the new sphere S' is the old one S , restricted to worlds where the floor gets mopped. The rule stated generally is

$$!\varphi: S \rightarrow S \cap |\varphi|,$$

or to formulate it as an identity:

$$(C) !\varphi(S) = S \cap |\varphi|.$$

The left-to-right inclusion here ($!\varphi(S) \subseteq S \cap |\varphi|$) follows from two extremely plausible assumptions:

(c1) commands shrink (i.e. don't expand) the sphere,

and

(c2) commands to φ make *all* $\sim\varphi$ -worlds impermissible

The right-to-left inclusion ($!\varphi(S) \supseteq S \cap |\varphi|$) follows from (c1) and a third plausible assumption

(c3) commands to φ make *only* $\sim\varphi$ -worlds impermissible.

All this is treated by Lewis as relatively undebatable, and nothing will be said against it here; it serves as the background to the problem to come.

That problem concerns permission. If commands go with intersection, the obvious first thought about permissions is that they would go with unions:

$$i\varphi: S \rightarrow S \cup |\varphi|,$$

or, the corresponding identity

$$(P) i\varphi(S) = S \cup |\varphi|.$$

The left to right inclusion ($i\varphi(S) \subseteq S \cup |\varphi|$) is hard to argue with; it follows from

(p1) permissions expand (i.e., do not shrink) the sphere,

and

(p2) permission to φ renders only φ -worlds permissible.

But the right to left direction requires along with (p1) the principle

(??) permission to φ renders *all* φ -worlds permissible.

And while it is hard to argue with

(p3) permission to φ renders *some* φ -worlds permissible,

(??) seems clearly wrong. Lewis explains why:

Suppose the Slave had been commanded to carry rocks every day of the week, but on Thursday the Master relents and says to the Slave, ‘The Slave does no work tomorrow’ . . . He has thereby permitted a holiday, but not just any possible sort of holiday . . . [not] a holiday that starts on Friday and goes on through Saturday, or a holiday spent guzzling in his wine cellar . . . (2000, p. 27)

So (??) allows in too much. (p3) on the other hand, although correct, can’t be the whole story. Not any old expanded sphere which contains φ -worlds will do, for the one whose sole φ -world has Slave staying on holiday through Saturday won’t do. So the situation is this:

Some worlds where the Slave does not work on Friday have been brought into permissibility, but not all of them. The Master has not said which ones. He did not need to; somehow, that is understood. (2000, p. 27)

If it’s understood, there must be a way we understand it: there must be a rule or principle of sphere-evolution that captures our shared implicit understanding of how permissions work.

Now we reach the problem of Lewis’s paper. What is that rule? Or to put it negatively, what exactly is wrong with a rule R that tells us that having been permitted to take Friday off, Slave can take that and other days off? Lewis looks at five answers.

(1) R lets in more worlds than necessary

Putting in a Saturday-off world enlarges the sphere more than necessary to allow Friday-off worlds. It’s a “gratuitous enlargement” in the sense of adding more worlds than necessary.

Lewis replies that any reasonable enlargement will be gratuitous in that sense, since the only non-gratuitous enlargements add in just a single world. This is fair enough, but it is not, I think, the “real” problem. If it were, then limiting ourselves to non-gratuitous (single-world) enlargements would address it. And it doesn’t; for we could pick as our single world a world where Slave takes Saturday off too.

(2) R lets in worlds more remote than necessary

Including Saturday-off worlds is a gratuitous enlargement in a *qualitative* sense. We should allow in only the *closest* worlds where the permitted action is done.

This, Lewis says, is too restrictive. Suppose Slave had previously been ordered to carry rocks around. Then he is forced to spend his vacation lifting weights! For weight-lifting worlds are closer to rock-carrying worlds than lying-around-at-the-beach worlds are to rock-carrying worlds.

One can put the problem like this. A permission should cleanly cancel relevant earlier commands; but on the present approach supposedly cancelled commands continue, from beyond the grave as it were, to exert an effect. The *clean cancellation requirement*, as I will call it, will come up again.

(3) R lets in worlds more impermissible than necessary

Allowing in Saturday-off worlds is a gratuitous enlargement, not in a qualitative but a prescriptive sense. We should put in the *least impermissible* worlds where the permitted action is done. Taking Friday and Saturday off was more impermissible than taking Friday off, so two-day-off worlds remain outside the sphere.

The objection Lewis offers is that this “solution” just restates, indeed aggravates, the problem: figuring out how comparative impermissibility evolves under the impact of commands and permissions is if anything harder than figuring out how straight permissibility does.

But there’s again a prior worry, I think—a version of the clean cancellation problem. Suppose Master first says not to eat any animals, then relents and permits Slave to eat lobster. Before lobster-eating was permitted, it was less impermissible to nibble on lobster than to eat a whole one. So afterward is it only permissible to nibble on lobster?

(4) R lets in worlds more disagreeable to Master than necessary

Allowing in Saturday-off worlds frustrates Master’s known or guessable purposes.

Lewis objects that either Slave knows Master’s purposes or he doesn’t. If he does, there’s no need for commands; he can work unsupervised. If he doesn’t, then the principle cannot be what’s guiding him.

Once again, there’s a prior worry. Let’s say Master has frequently ordered Slave to carry rocks up the hill. Presumably she did this because she wants the rocks up the hill. But the Friday-off worlds that best serve the purpose of getting them up the hill are ones where Slave invites his friends to a game of “let’s see who can carry more rocks up the hill”. This is again a version of the clean cancellation problem.

(5) R lets in worlds violating more commands than necessary

This takes a bit more explanation. It's a given that Master doesn't issue commands and permissions unless she needs to. She doesn't issue the command to φ if it is already impermissible for Slave not to φ ; and she doesn't issue permission to φ if it is already permissible for Slave to φ . In particular, then, Master would not have permitted Slave to take Friday off unless taking Friday off would otherwise have been an act of disobedience, an act in violation of some explicit or understood command. So, proposal: the effect of permitting φ should be to invalidate any commands that forbid φ -ing—that are inconsistent with φ —while leaving other commands in place. The problem with an update rule that lets Saturday-off worlds into the sphere is that it invalidates more commands than necessary. To make Slave's taking Friday off permissible, it's enough to invalidate the work-*Friday* command; the work-Saturday command doesn't care if Slave takes Friday off, so it should be left in place.

Call this update rule the *remainder* rule because it defines S^+ as the set of worlds satisfying the commands that remain when the φ -inconsistent commands are knocked out. Lewis doesn't like this rule either; here is why. Clearly, to apply the rule, we need there to be a list of commands ψ_1, \dots, ψ_k such that a world is permissible iff it complies with all of them, that is,

$$S = |\psi_1| \cap |\psi_2| \cap \dots \cap |\psi_k|.$$

For the way the rule works is we delete from this list all the ψ_i s inconsistent with φ , and let the commands that remain define S^+ . If the φ -incompatible commands are $\psi_{j+1}, \psi_{j+2}, \dots, \psi_k$, the new sphere is

$$S^+ = |\psi_1| \cap |\psi_2| \cap \dots \cap |\psi_j|.$$

Where is the initial set of commands supposed to come from, though, the one we thin out to arrive at the reduced command-set that defines the S^+ -worlds? It would be one thing if “ $i\varphi$ ” were the first permission uttered; for then Master's earlier utterances were all commands and we can let the ψ_i s be those commands. Ordinarily, though, “ $i\varphi$ ” is preceded by commands *and other permissions*. One could try considering just the commands that have already been given, ignoring the permissions, but these will not define the current sphere of permissibility, because the update effects of the permissions will have been ignored.

It seems, then, that we are driven to *contriving*, reverse-engineering if you like, a package of commands that define the current sphere. Unfortunately the relation between S and commands defining S is one-many; lots of packages will issue in the same sphere of permissibility. How does Slave know which package to use? It makes a difference, because the effect on S of permitting φ varies enormously with our choice of implicit commands ψ_i .

Suppose that the current sphere S = the worlds where Slave works all day, every day from Monday to Sunday; and that we arrived at that sphere by a series of enlargements and contractions that offer no clues to what the right ψ_i s, the right implicit commands, are. Slave *might* think that initially, before he is given Friday off, the commands in effect are

- ψ_1 : Slave carries rocks on Monday.
- ψ_2 : Slave carries rocks on Tuesday.
- ψ_3 : Slave carries rocks on Wednesday.
- ψ_4 : Slave carries rocks on Thursday.
- ψ_5 : Slave carries rocks on Friday.
- ψ_6 : Slave carries rocks on Saturday.
- ψ_7 : Slave carries rocks on Sunday.

The one command here inconsistent with “Slave takes Friday off” is “Slave carries rocks on Friday”. Suspending that one command leaves the commands to work other days still in place. Clearly on this way of doing it, Slave has *not* been permitted to take other days off, which was the desired result. But Slave might also think that the implicit commands are

- χ_1 : Slave carries rocks on weekdays.
- χ_2 : Slave carries rocks on the weekend.

Now the φ -inconsistent rule, the one to be cancelled on the present hypothesis, is “Slave carries rocks weekdays”. But then the sphere of permissibility expands to include all worlds where Slave works on the weekend. And that seems crazy. Master meant to give Slave Friday off, but not Monday–Thursday as well.

Lewis’s objection in a nutshell is that the implicit commands are too unconstrained for the remainder rule to be of any use. He may be right in the end. I wonder, though, whether there are constraints he is missing—constraints that don’t come into view until you raise his sort of problem in the starkest possible terms. Let’s look then at the most extreme cases of badly chosen implicit commands. At the one extreme we have commands *each* of which is inconsistent with φ ; at the other we have ones *none* of which is inconsistent with φ . An example of each-inconsistent is

- θ_1 : Slave carries rocks every morning of the week.
- θ_2 : Slave carries rocks every afternoon of the week.

Neither of these is compatible with Slave taking Friday off. Cancelling the φ -inconsistent commands, then, is cancelling all commands whatsoever. If all commands are canceled, then everything is permitted. Master wanted to let Slave take Friday off, but winds up giving Slave his freedom.

Now consider commands none of which individually requires Slave to work Friday, but whose *joint* effect is to require Slave to work every day of the week. For instance,

- σ_1 : Slave carries rocks every morning if any afternoon.
- σ_2 : Slave carries rocks every afternoon if every morning.
- σ_3 : Slave carries rocks some afternoon.

σ_1 allows Slave to take Friday off, provided he never carries rocks in the afternoon. σ_2 allows him to take Friday off, provided he omits to carry rocks some morning. σ_3 allows him to take Friday off, provided he carries rocks some afternoon. Each of the σ_i s is consistent with φ , so none of them is canceled on the present rule. But then the sphere of impermissibility never changes. Master tried to give Slave permission to take Friday off, but it turns out he still has to work on Friday.

What can we conclude from this? The remainder rule—the one that says to cancel all and only pre-existing commands that forbid φ —can give *very silly results*. But we can make that work in our favor, by using the silliness of these results as a way of tightening up the rule. Call a command-list *reasonable* if running it through the remainder rule yields an expansion satisfying (p1)–(p3) above.

- (p1) permissions expand (do not shrink) the sphere, and
- (p2) permission to φ renders *only* φ -worlds permissible.
- (p3) permission to φ renders *some* φ -worlds permissible.

It is not hard to establish the following (proof in Appendix to the chapter):

FACT: If S is defined by a reasonable list of commands, then $S = |\sim\varphi| \cap |\psi|$ for some ψ . Equivalently, any reasonable command list is of the form (up to equivalence) “You must not φ ,” “You must ψ .”

This transforms the problem in a helpful way. Before we had one equation in several unknowns (corresponding to the several choices of implicit commands ψ_i). Now we have something very like one equation in one unknown. For we know what S is; that’s just the present, pre-permission-to- φ , sphere of permissibility. And we know what $|\sim\varphi|$ is; that’s just the worlds where the permitted behavior does not occur. The one unknown here is $|\psi|$, that is, the new sphere of permissibility S^+ .

So, to review. Whenever a permission to φ is issued, it’s as though the initial command list had consisted of two commands:

- first, one saying (precisely) *do not* φ .
- second, a command that allows φ -ing,

Our job as sphere-redrawers is to throw out the *do not* φ command and form the set of worlds allowed by the command that remains. This is nothing like an algorithm, because there is more than one way of choosing the command that remains. (There are many sets whose intersection with $|\sim\varphi|$ is S .) But it is instructive nevertheless.

One way it helps is by showing how to conceive the task diagrammatically. We are given the $|\varphi|$ -region—that’s the worlds where Slave takes Friday off, as he is permitted to do. We are given the $|\sim\varphi|$ -region—that’s the worlds where Slave works Friday, in accord with his pre-permission obligations. We are given the S -region—that’s the set of initially permissible worlds where Slave works all week (Monday–Sunday). Our job is to extrapolate the S -region beyond the bounds imposed by the $|\sim\varphi|$ -region, thus arriving at the set $|\psi|$ of worlds that are permissible after Master cancels the command to work Friday.

Three observations about this diagram. *First*, the diagram helps us to see what it means to say that S^+ “solves the equation” $S = |\sim\varphi| \cap S^+$. It means that S is the part of S^+ where the newly permitted behavior doesn’t occur, or to run it the other way around, S^+ is the result of extending the S -region into the region where the newly permitted behavior does occur.

A *second* thing the diagram helps us to see is that the extrapolation approach is no panacea. There is not going to be just a single way of extending S into

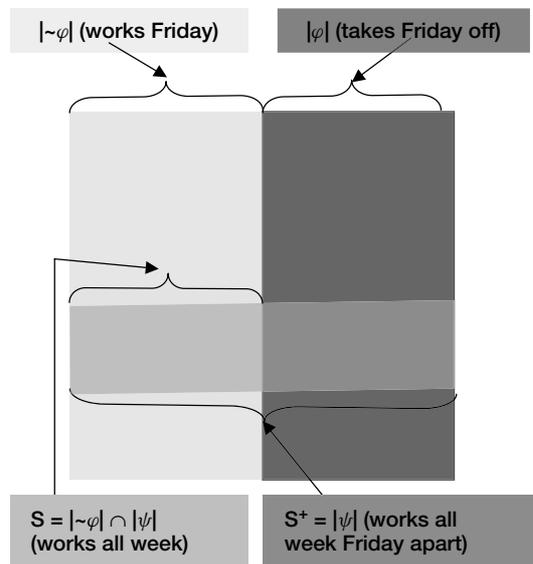


Fig. 11.1.

the φ -region, the region where Slave takes Friday off. One could in principle let S^+ suddenly triple in width as it crosses the φ -border, so that Slave is allowed to take the weekend off too. And various other improper extrapolations can be imagined.

A *third* thing that the diagram helps us to see, however, is that some ways of extrapolating are more natural than others. Extending S “directly” into the φ -region, ignoring the $\varphi/\sim\varphi$ boundary, is better than taking notice of that boundary. This is the geometrical upshot of our remarks above about clean cancellation. To the extent that an extrapolation takes account of the $\varphi/\sim\varphi$ boundary, the supposedly canceled prohibition on φ -ing is still exerting some influence from beyond the grave. Another way to put it is that we want our remainder command ψ , aka S^+ , to be *free* of any taint of the canceled command $\sim\varphi$.

Our problem boils down to this: what could it mean for one proposition (in this case, $|\psi|$) to be “free” of another ($|\sim\varphi|$), in the sense of indifferent to whether that other proposition is true or false. The following answer looks plausible:

A is B-free iff
in worlds where A is true, it is true for reasons compatible with B’s falsity,
and in worlds where A is false, it is false for reasons compatible with B’s truth.¹²

This notion of “true for the same reason” will have to be left at an intuitive level, but it basically means true with the same truth-maker—and here one wheels in an appropriate theory of truth-makers. The proposed update rule for permissions is this:

(URi)
Suppose that S is the present sphere of permissibility and that φ holds in no S -worlds. $S^+ = S + i\varphi$ iff four conditions are met:
Difference $S^+ - S$ is non-empty
Equality $S = |\sim\varphi| \cap S^+$
Freedom S^+ is $|\sim\varphi|$ -free
Goodness Other satisfiers of D, E, and F are “less good”.

For our purposes today, the fourth condition is just to ensure that “ $+ i\varphi$ ” is a function.¹³ It is a non-trivial question what makes one otherwise qualified contender for the role of S^+ better than another; it could be a matter of relative naturalness, or logical strength, or something else again. That does not have to be decided here, however; all the work will be done by conditions D, E, and F.

Suppose that S , our initial sphere of possibility, is the proposition that Slave works all week, and $|\sim\varphi|$ is the proposition that Slave works on Friday. How

¹² Perhaps we should say, “in all not too remote worlds where A is true/false . . .”.

¹³ Albeit perhaps a partial function.

good a candidate is the proposition that Slave works Monday–Thursday and the weekend for the role of what’s permitted after Slave is given Friday off?

Is the difference condition met? Yes—not all worlds where Slave works Monday–Thursday and the weekend are worlds where Slave works Monday thru Sunday.

Is the equality condition met? Yes, for the work-Monday-to-Sunday worlds are the intersection of the work-Monday-to-Thursday-and-the-weekend worlds with the work-Friday worlds.

Is the freedom condition met? In worlds where S^+ is true, it is made true by Slave’s rock-carrying activities Monday–Thursday and the weekend. Those activities, whatever exactly they may be, are compatible with Slave taking Friday off. In worlds where S^+ is false, it is made false by Slave’s non-rock-carrying activities on some day or days other than Friday. Those activities, whatever they may be, are compatible with Slave carrying rocks on Thursday. “Slave works Monday–Thursday and the weekend” is thus free of any taint of “Slave works Friday”. That and the equality fact already mentioned are what makes “Slave works Monday–Thursday and the weekend” a good candidate for what is left over when we subtract Slave’s working Friday from his working every day of the week.

Now let’s consider some other hypotheses about what is ultimately required of Slave when we first command him to work Monday–Sunday and then relent, allowing him to take Friday off.

Harsh hypothesis: Slave still has to work Monday–Sunday, exactly as before.

Lenient hypothesis: Slave only has to work Monday–Thursday; after his Friday holiday he can stay off through the weekend.

Strange hypothesis: Slave has to work Monday–Thursday, the weekend, and Master’s birthday.

The harsh hypothesis—Slave still has to work every day—falls afoul of the Difference condition, which says permission to φ , issued in a context where φ was previously impermissible, must result in the addition of at least one φ -world to the previous sphere. Obviously no Friday-off worlds have been added if S^+ is (like S) the set of work-Monday-to-Sunday worlds.

The lenient hypothesis—Slave only has to work Monday–Thursday—falls afoul of the Equality condition. The work-everyday worlds are not identical to but a proper subset of S^+ (the work-Monday-to-Thursday worlds) intersected with $|\sim\varphi|$ (the work-Friday worlds).

The strange hypothesis—Slave works Monday–Thursday, the weekend, *and* Master’s birthday—falls afoul of the Freedom condition, by which I mean that it is not free of any taint of the canceled work-Friday command.

A is B-free, recall, iff A is true, in an A-world w , for reasons compatible with B’s falsity, and false, in a $\sim A$ -world w , for reasons compatible with B’s

truth. But consider a world w where Master's birthday falls on a Friday, and Slave works Monday–Thursday, the weekend, and Master's birthday, as the strange hypothesis requires him to. The hypothesis's truth-maker involves in part that Slave works Friday, given that Friday is Master's birthday. And this is not compatible with $\sim\varphi$ being false, in other words with Slave's taking Friday off. Alternatively consider worlds where Master's birthday falls on a Friday, and the strange hypothesis *doesn't* hold, because Slave takes Friday off. Is the falsity-maker here compatible with $\sim\varphi$'s truth, that is, with Slave's working Friday? Clearly not. So our update rule predicts that Slave is not required to work Monday–Thursday, the weekend, *and* Master's birthday, *even where that birthday is not on Friday*, because in some not-too-distant worlds the command *would* be obeyed by working on Friday, or disobeyed by taking Friday off; and what Slave does on Friday should not make him either obedient or disobedient, once he has been freed of all Friday-related obligations.

That is my first-pass story about how to understand the semantic effect of permissions. There are gaps in it, to be sure. One large gap traces back to how Lewis sets the problem up. Lewis assumes that permission to φ “makes a difference” if, *but only if*, the sphere has to be enlarged for there to be permissible worlds in which φ . If there were any permissible φ -worlds beforehand, then the permission accomplishes exactly nothing.

This may seem only reasonable. The first day of camp, Counselor says: “It is forbidden to climb trees.” The director then whispers something in her ear, and she adds, “except you may climb trees to rescue a kitten”. Here “you may climb trees to rescue a kitten” deletes something from the list of what was ruled off limits by “It is forbidden to climb trees”. It might seem at first that all permission statements are like that. What would be the point of allowing something which had never been forbidden in the first place?

And yet we do it all the time. I had Counselor first forbidding tree climbing and then relenting a bit: “It is OK to climb trees to rescue a kitten.” But she might just as well have granted the permission first. There would be a clear point to doing this. If Counselor says nothing, then campers don't know if tree climbing is forbidden; perhaps she hasn't got around to announcing it yet. If she says tree climbing after kittens is OK, that tells them she is rejecting the opportunity to announce it.

There are also intermediate cases where a permission “softens” earlier commands, even though the earlier commands did not strictly rule out behavior of the kind now permitted. On day one, Counselor says, “You must never climb trees.” On day two, she says, “You may do whatever you like on your birthday.” Lewis would, it seems, have to say that the permission on day two leaves the command on day one entirely in place, since there are plenty of worlds where

you never climb trees and still do whatever you like on your birthday, for instance, worlds where you have no desire to climb trees. Intuitively, though, the day one command is weakened. More on this in a moment.

Now I want to explore the epistemic analogue of Lewis’s Master–Slave game. Here is how I understand the new game to work.

- (1) The players this time are Teacher and Student, and the sphere of permissibility becomes the sphere of believability.
- (2) The old game had Slave constantly adjusting his plans to fit with changes in what was permissible; the new one has Student constantly adjusting his theory to fit with changes in what is believable.
- (3) It contracted the sphere of permissibility when Master said, “Do ψ ”; the sphere expanded when Master said, “You may do φ ”. Likewise it contracts the sphere of believability when Teacher says “ ψ is so”; the sphere expands when Teacher says “ φ may be so”.
- (4) There was no great mystery about the *kind* of contraction brought on by “Do ψ ”; one simply rejected as impermissible worlds where ψ failed. Similarly there is no great mystery about the kind of contraction brought on by “ ψ is so”; worlds where ψ fails are rejected as unbelievable.
- (5) It was initially mysterious how “You may do φ ” enlarged the sphere of permissibility. Similarly it is mysterious to begin with how “ φ may be so” enlarges the sphere of believability.

Let’s continue the pretense that “ φ may be so” has no effect on a sphere of believability that contains φ -worlds; it’s only when all believable worlds are $\sim\varphi$ that we get an expansion. The question is, what expansion do we get? I propose that the update rule is pretty much as before.

(UR \diamond)

Suppose that S is the present sphere of believability and that φ holds in no S -worlds. S^+ is $S + \diamond\varphi$ iff four conditions are met:

- Difference* $S^+ - S$ is non-empty
- Equality* $S = |\sim\varphi| \cap S^+$
- Freedom* S^+ is $|\sim\varphi|$ -free
- Goodness* Other satisfiers of D , E , and F are “less good”.

Once again, the point of the fourth condition is to ensure that “+ $\diamond\varphi$ ” is a function.¹⁴ I am not sure what makes one otherwise qualified contender for the role of S^+ better than another. Naturalness plays a role, presumably, and perhaps

¹⁴ Albeit perhaps a partial function.

also logical strength. This does not have to be decided here, however; all the work will be done by conditions D, E, and F.

Imagine that Teacher starts by saying it will rain all week, meaning Monday–Sunday. She thereby banishes from the sphere of believability all worlds where this doesn’t happen, where the rain lets up on one or more days. When Teacher learns that her evidence as regards Friday was shaky, she says, “Hold on, it might not rain Friday after all.” Which worlds is Student to put back into the sphere of believability? Or to put it in more intuitive terms, what remains of Teacher’s original prediction of rain all week, once she has conceded it might not rain Friday?

One hypothesis, the strong hypothesis call it, has Teacher still predicting that it rains every day. That is not allowed by our update rule, however, for it violates Difference. Difference says that new worlds *have* to be added to accommodate a previously unbelievable hypothesis. And it is clear from Equality that the added worlds have to be φ -worlds; for if a $\sim\varphi$ -world w is added, then S is a proper subset of $|\sim\varphi| \cap S^+$, not equal to it as required by Equality. Teacher’s announcement that it might not rain Friday forces the addition of at least one dry-Friday world to the sphere of believability, so the strong hypothesis is mistaken.

A second hypothesis, the weak hypothesis call it, has Teacher now predicting only that it rains Monday–Thursday. That too is not allowed by our update rule. For suppose S^+ is the set of worlds where it rains Monday–Thursday. Then S^+ intersected with $|\sim\varphi|$ (the set of wet-Friday worlds) contains worlds where Saturday and Sunday are dry. There are no such worlds in S , however; before Teacher allows it might be dry Friday, all believable worlds have it raining Monday–Sunday. It follows that S^+ intersected with $|\sim\varphi|$ contains worlds outside of S , which is contrary to Equality. So the weak hypothesis about what remains of the Teacher’s original prediction is mistaken too.

A third, intermediate, hypothesis has Teacher now predicting that it rains Monday–Thursday, the weekend, and on Teacher’s birthday. The problem with this is that “It rains on Teacher’s birthday” is not free of any taint of the canceled wet-Friday assertion; for in worlds where Teacher’s birthday falls on Friday, either “It rains on Teacher’s birthday” is true for reasons incompatible with the falsity of “It rains Friday”, or it is false for reasons incompatible with the truth of “It rains Friday”. Our update rule objects to “It will rain Monday–Thursday, the weekend, and Teacher’s birthday,” *even if her birthday is not on Friday*, because in some worlds it *is* on Friday. In those worlds the prediction *is* verified (falsified) by Friday’s lack of rain; and Friday’s lack of rain ought to be irrelevant given that Teacher freely admits there might be no rain on Friday.

I want to return now to a limitation of the Lewis game noted earlier. Lewis stipulates that permission to φ has no impact unless φ was antecedently forbidden. That makes nonsense both of out-of-the-blue, discourse-initial, permissions

and permissions that weaken earlier commands that didn't strictly forbid the now-permitted behavior. The same points apply to our epistemic analogue of the Lewis game. We have been assuming that "might φ " has no impact unless φ was antecedently denied. But then what is going on in this conversation?

A: Where is Bob?

B: Hmmm, don't know for sure, but he might be in his office.

A: *I never said he wasn't!

Or this one?

A: Bob will be at the office tomorrow.

B: Not so fast, he might still have the flu tomorrow.

A: *That's compatible!

Lewis confronts what might be considered the dual of this difficulty in "A Problem About Permission". Having laid it down that commands shrink the sphere of permissibility, he remarks that

One sort of commanding may seem to require special treatment: commanding the impermissible. Suppose that $|\varphi|$ contains no worlds that are . . . permissible . . . The Master may nevertheless wish to command . . . that φ . . . Having commanded at dawn that the Slave devote his energies all day to carrying rocks, the Master may decide at noon that it would be better to have the Slave spend the afternoon on some lighter or more urgent task. If the master simply commands . . . that φ , then no world . . . remains permissible; the Slave, through no fault of his own, has no way to play his part by trying to see to it that the world remains permissible . . . Should we therefore say that in this case the sphere evolves not by intersection but in some more complicated way? (2000, p. 27)

He notes a possible fix: whenever φ is impermissible, "a command that φ is deemed to be preceded by a tacit permission that φ , and the sphere of permissibility evolves accordingly" (2000, p. 27). Our present concern can be put in similar language:

One sort of permitting may seem to require special treatment: permitting the not impermissible. Suppose that the sphere of permissibility contains φ -worlds. The Master may nevertheless wish to permit that φ . Having at dawn permitted the Slave to take the day off, the Master may decide at noon that the Slave should be permitted to visit his mother this week. If the Master simply permits the Slave to visit his mother this week, then no additional worlds become permissible; for there are already permissible worlds where the Slave visits his mother, namely worlds where the Slave visits his mother today. Should we therefore say that in this case the sphere evolves not by the remainder rule but in some more complicated way?

I suggest we can avoid saying this by a maneuver similar to Lewis's: whenever φ is already permissible, permission to φ is deemed to be preceded by a tacit command not to φ , with the sphere of permissibility evolving accordingly.

Likewise whenever φ is already believable, “it might be that φ ” is imagined to be in response to the unspoken assertion that $\sim\varphi$.

How much justice does this kind of maneuver do to our feeling of still *conveying* something when we permit the not previously permissible, or suggest that things *might* be a way that no one had ever said they weren’t?

The first thing to notice is that, just as permitting and then immediately commanding that φ can (even by our existing rules) change the sphere of permissibility, forbidding and then permitting φ can change the sphere of permissibility too. Mathematically speaking there is no reason whatever to expect that $S^{-+} = i\varphi (!\sim\varphi(S))$ will just be S again. Indeed there is reason to expect it often won’t. We know by Freedom that

$$S^{-+} \text{ is } |\sim\varphi|\text{-free.}$$

It follows that whenever S (which is arbitrary, recall) is *not* $|\sim\varphi|\text{-free}$, S^{-+} is not S , which is just to say that the operation of forbidding and then permitting φ will have non-trivial effects. Example: we saw above that $|\text{Slave works on Master’s birthday}|$ is not free of $|\text{Slave works on Friday}|$. So if we let S be the first of these and $\sim\varphi$ be the second (so φ says that Slave takes Friday off), we should have a case where $S^{-+} \neq S$.

Let’s try it. Master first says that Slave is to work on her (Master’s) birthday. That gives us the desired S . Then Master further commands that Slave is to work on Friday. That gives us $S^- = !\sim\varphi(S) =$ the worlds where Slave works Friday and Master’s birthday. Then Master permits Slave not to work on Friday. S^{-+} can’t be the worlds where Slave works on Master’s birthday again, because that is not free of Slave’s working on Friday, which he has just been permitted not to do. A better because more Friday-free candidate for S^{-+} is the set of worlds where Slave works on Master’s birthday unless Master’s birthday falls on a Friday. So permitting Slave to take Friday off in a context where Slave was required only to work on Master’s birthday has a non-trivial effect on the sphere of permissibility.

Like remarks apply to asserting that φ and then immediately allowing it might be that $\sim\varphi$. Suppose I’ve asserted that it will rain on my birthday. Allowing it might not rain on Friday has the effect, I’m suggesting, of asserting it will rain on Friday and then taking it back. Once again this does not leave everything as it was. The prediction that remains is that it will rain on my birthday, provided that my birthday doesn’t fall on a Friday.

I have argued that forbidding and then immediately permitting φ can change the sphere of permissibility, and also that asserting φ and then immediately allowing that maybe φ is not the case can change the sphere of believability.

But there are also cases where permitting what I’ve just forbidden (admitting that a previous assertion might be wrong) leaves the sphere just as it was. An example might be this. Nothing has been said about Bob’s location, but I know

you want to find him. What is accomplished by saying, “He might be in his office,” when no one has suggested otherwise? Likewise what is accomplished by announcing out of the blue that it is permitted to climb trees in order to rescue kittens?

It seems to me these things are not so mysterious, once we distinguish what *has* been forbidden, in the sense that the command has been given, and what *is* forbidden, in the sense that it’s against the rules but Counselor may not have got around to announcing it yet. The children may know when the first permission is given that nothing *has been* forbidden, but they have no idea what might or might not *be* forbidden in the sense of being off limits or against the rules. When they hear that tree climbing after kittens is permitted, they learn an upper bound on what *is* forbidden, namely that it doesn’t include tree climbing after kittens. This is not because Counselor has *said* climbing after kittens is not forbidden; she has the ability to forbid but not, as we’re imagining the game, the ability to comment on the extent of the forbidden. What the counselor has done is “shown” that climbing after kittens is not forbidden by staging a confrontation with an imagined off-screen forbiddler, and *canceling* that imagined person’s decree.

Something similar is going on when I say to someone looking for Bob that he might be in his office. The distinction we need this time is between what has been asserted, and what is understood to be so even if no one has got around to announcing it yet. Before I spoke, my friend might have been wondering whether an assertion that Bob was not in his office was in the cards. I satisfy her curiosity not by *saying* that an assertion to that effect is not in the cards; my subject matter is Bob and his office, not assertions about them. I satisfy my friend’s curiosity by *showing* that an assertion to that effect is not in the cards, by staging a confrontation with someone imagined to have made the assertion, and undoing what they are imagined to have done.

Let’s return now to some of the problems raised at the outset, starting with problems for the standard semantics (SS). One problem was that SS gave “might φ ” the wrong subject matter. “Bob might be in his office” seems intuitively to be about whatever “Bob is in his office” is about. Neither concerns the speaker or the extent of her knowledge. The present view construes “might φ ” as a device for retracting or canceling an assertion of $\sim\varphi$. If φ has the same subject matter negated as unnegated, and $\sim\varphi$ has the same subject matter retracted as asserted, then “might up” comes out with the same subject matter as φ .

A second worry was that the truth-conditions assigned by SS, in its Moorean form at least, were too weak. If “might φ ” says only that *my* information doesn’t rule φ out, why do I accept correction by observers with information that I didn’t possess?¹⁵ It is indeed puzzling why I would accept correction, if that means agreeing I have misstated the epistemic facts. But suppose that “might

¹⁵ See however the “Mastermind” example in von Fintel and Gillies (2008).

φ ” is not a statement of fact. Suppose it is a “cancellation order”, an attempt to undo or reverse the assertion that $\sim\varphi$, an insistence that $\sim\varphi$ not be part of the common ground. If *that* is what the observer is taking issue with, then her point can be understood as follows: however well intentioned, the cancellation order was unfortunate. Given that somebody knew that φ , it would have been better not to block φ ’s addition to the common ground.

Now, as we discussed, the standard semanticist’s response to the “too weak” objection is to make the truth-conditions stronger: φ should be consistent not only with *my* information, but all *pertinent* information—where the test of pertinence must presumably be that the speaker concedes that if that information really does obtain, then he was in error. The third worry was that these revised truth-conditions are too strong. If “might φ ” is false when φ is ruled out by pertinent facts, then speakers should restrain themselves except when such facts are known not to obtain. I shouldn’t say that Bob might be in his office, if there’s a chance that Bob has unbeknownst to me been seen elsewhere. Clearly, though, speakers do *not* restrain themselves in this way. (If they did, “might”-claims would hardly ever get made.) What can the cancellation theory say about these sorts of cases? How much restraint to expect depends on what speakers are afraid of. If I am afraid that my claim might be *false*, then I should hold back until I have tracked down all pertinent facts. But what if I am concerned only that my claim will turn out to have been ill-advised, or counterproductive, given the purposes of the conversation? The claim is counterproductive just to the extent that it pre-empts or cancels better-informed assertions of $\sim\varphi$ by others. And I *do* restrain myself from saying “might φ ” when there is a danger of this. If someone announces over the phone that they are looking at Bob right now, I am hardly likely to tell them that Bob might be in his office.

The fourth problem was that SS is too epistemic. It reckons “I might vote for Kucinich and I might not” false unless I am genuinely undecided how I am going to vote. The present theory can say that I am showing my audience, by example, as it were, that no assertion is to be expected on the topic of how I am going to vote; I do it by giving myself an opportunity to make that sort of assertion and visibly passing it up.

The fifth problem was that “Bob might be in his office” does not entail “Bob might be in his office or in an opium den,” as SS would lead you to expect. “Bob might be in his office or in an opium den” seems to imply both that he might be in his office and that he might be in an opium den. The analogous phenomenon with permission is perhaps better known. Suppose you are hungry and I tell you: You may have a piece of cake or a piece of pie. You reach for the pie and I snatch it away. What gave you the idea that *that* was a permissible disjunct?

A “stronger statement” in the context of the cancellation theory is a statement that cancels more; so it is enough to show that whereas $\diamond\varphi$ cancels only $\sim\varphi$, $\diamond(\varphi$

or ψ) cancels $\sim\varphi$ and $\sim\psi$ both. Let S be a sphere of believability that implies $\sim\varphi$ and implies $\sim\psi$. The update rule for “might” tells us that $S + \Diamond(\varphi \vee \psi)$ is a superset S^+ of S that is free of $\sim(\varphi \vee \psi)$. By the definition of freedom, S^+ is free of $\sim(\varphi \vee \psi)$ only if

(#) S^+ is false, when it is, for reasons compatible with the truth of $\sim(\varphi \vee \psi)$, that is, for reasons compatible with the combined falsity of φ and ψ .

We have, then, that facts making S^+ false are compatible with the joint falsity of φ and ψ . Facts making either φ or ψ true are not compatible with the joint falsity of φ and ψ . So facts making either of φ or ψ true do not make S^+ false. If S^+ implied $\sim\varphi$, however, then facts making φ true presumably *would* make S^+ false; and there would be a similar issue if S^+ implied $\sim\psi$.¹⁶ Therefore S^+ does not imply $\sim\varphi$ and it does not imply $\sim\psi$; which, bearing in mind that S itself does imply $\sim\varphi$ and $\sim\psi$, means that the effect of $\Diamond(\varphi \vee \psi)$ is to cancel $\sim\varphi$ and $\sim\psi$ both. This accounts for the feeling that “It might be that φ or ψ ” implies both “It might be that φ ” and “It might be that ψ ”.

The sixth and final problem we raised for SS is that it has trouble explaining why “ φ & it might be that $\sim\varphi$ ” is not coherently supposable, e.g., in the antecedent of a conditional, when no such problem arises with “ φ & I/we don’t know that φ ”. The problem with “ φ & it might be that $\sim\varphi$ ” is not that no *world* can answer both to the specification that φ and the specification that $\Diamond\sim\varphi$; it’s that no *world-specification* can both demand that φ (as the first conjunct requires) and not demand that φ (as the second conjunct requires).

Update semantics, we said, left it un- or under-explained why allowing that Ringo might not go to the party leaves intact the information that John, Paul, and George will be there. The present theory says that it is only prior assertions bound up with (= not free of) Ringo’s being at the party that get canceled. “John, Paul, and George will attend” is free of “Ringo will attend” in that facts making the first true (false) do not conflict with facts making the second false (true). This is just one more instance of Lewis’s problem of clean cancellation: prior assertions that entail the falsity of what we learned might be true are canceled, while prior assertions that do not entail the falsity what we learned might be true are left in place.

Our second worry about update semantics was this. It tells us that $\Diamond\varphi$ uttered in information state S has no effect unless S and φ are inconsistent. Suppose that $S = \{\sim\chi\}$ and $\varphi = \Diamond(\chi \vee \psi)$. (χ or ψ) is consistent with $\sim\chi$, so $S + \varphi$ should be S again. But we know that $\Diamond(\chi \vee \psi)$ generally entails $\Diamond\chi$ and $\Diamond\psi$. And to be told that it might be that χ presumably cancels any prior information we had to

¹⁶ The inference here does not go through in every case. For φ ’s truth-makers need not always be falsity-makers for each χ that implies $\sim\varphi$. Let $\chi =$ “There’s a French king in my pocket” and $\varphi =$ “France has no king”. φ is false because France is a republic but χ ’s falsity is due to the fact that my pocket is empty.

the effect that $\sim\chi$. Here then is a case where $S + \varphi$ is a proper subset of S , even though φ is consistent with S .

How does the cancellation theory deal with this case? $S + \diamond\varphi$ is the $\sim\varphi$ -free part of S . From the fact that φ is consistent with S , that is, S does not *imply* $\sim\varphi$, we cannot infer that S is absolutely *free* of $\sim\varphi$. An example has just been mentioned: (χ or ψ) is consistent with $\sim\chi$ and $\sim\psi$ —it does not imply χ or ψ —but a disjunction can hardly be considered free of its disjuncts. *Non-implication of χ* is a much weaker property than *freedom from χ* .

This completes my explanation and defense of the cancellation theory. I am aware that important topics have been left undiscussed: the Frege–Geach problem, for instance, and the problem of how deontic/epistemic peers discuss what to do/think.¹⁷ I would like to take the opportunity in closing to cancel any would-be assertions to the effect that these topics will not be discussed in future work.

APPENDIX

FACT: Any S defined by a reasonable command list = $|\sim\varphi| \cap |\psi|$ for some ψ .

PROOF: From (p3), which says that the right expansion should bring in at least one φ -world, we conclude that any package of commands each of whose members is consistent with φ is unreasonable. Such a package fails to enlarge the sphere of permissibility, as it has to be enlarged to make room for φ -worlds. From (p2), which says that the right expansion should bring in *only* φ -worlds, we conclude that any package of commands none of whose members is consistent with φ is unreasonable. The reason is that that kind of package expands the sphere of possibility to include *every* world, and we know by (p2) that permission to φ should bring in only some worlds, only φ -worlds. So, any reasonable package of commands $\langle \psi_i \rangle$ has members consistent with φ and members inconsistent with φ . Let's use χ for the conjunction of all ψ_i s *inconsistent* with φ , and ψ for the conjunction all ψ_i s individually *consistent* with φ . Then

¹⁷ The games I discuss are unrealistic in that one of the players is entirely in charge. A more realistic version of the Master–Slave game would have the players working out what to do together, with an eye to some shared goal. Each makes proposals about how to proceed which may or may not be accepted by the other. When a “let’s do this” proposal is accepted, the sphere of permissibility shrinks; when a “let’s allow ourselves not to do that” proposal is accepted, the sphere grows. A more realistic version of the Teacher–Student game would have two parties working out what to think, or assume, together. Assertions are proposed lower bounds on an evolving body of shared assumptions. If accepted, they shrink the sphere of believability, or at least stand in the way of certain imaginable expansions. Might-statements are proposed upper bounds on our shared assumptions; if accepted they expand the sphere, or else stand in the way of imaginable proposals to shrink it. The Master–Slave game thus modified brings in a degree of mind-to-world direction of fit, since certain proposals about what to do are objectively unwise given the players’ goals. The Teacher–Student game thus modified has a world-to-mind element; the players are haggling, wikipedia-style, about what goes into the body of shared assumptions—what can be taken for present purposes as settled fact. (Thanks here to Sally Haslanger and Bob Hale.)

S = the set of $(\chi \wedge \psi)$ -worlds
 S^+ = the set of ψ -worlds
 $S^+ - S$ = the set of $(\psi \wedge \sim\chi)$ -worlds

Again, $\langle \psi_i \rangle$ is reasonable only if

\exists some worlds in $S^+ - S$ are φ -worlds (from (p3))
 \forall all worlds in $S^+ - S$ are φ -worlds (from (p2))

From \exists it follows that ψ is consistent with φ . For: suppose not. Then $S^+ =$ the set of ψ -worlds does not contain any φ -worlds. But \exists implies that S^+ does contain φ -worlds, since $S^+ - S$ contains them. A more interesting result follows from \forall : for all S^+ -worlds w , χ holds in w iff φ does not hold in w . The “only if” direction is easy since each of χ 's conjuncts is by definition inconsistent with φ . For the “if” direction, suppose contrapositively that χ does not hold in w . w cannot be an S -world because S -worlds have to satisfy all the ψ_i s. But then w is in $S^+ - S$. And according to \forall , every world in that set satisfies φ . So, the implicit commands suitable to serve as backdrop to a permission to φ must be divisible into two parts: first, $\chi \approx \sim\varphi$ = the part that forbids φ -ing, second, ψ = the part that allows φ -ing. QED