I Believe I Can φ^*

Matthew Aaron Mandelkern, Ginger Schultheis and David Boylan

Massachusetts Institute of Technology, Cambridge, Massachusetts, U.S.A.
`mandelk@mit.edu, vks@mit.edu, dboylan@mit.edu`

Abstract

We propose a new analysis of ability modals. After briefly criticizing extant approaches, we turn our attention to the venerable but vexed conditional analysis of ability ascriptions. We give an account that builds on the conditional analysis, but avoids its weaknesses by incorporating a layer of quantification over a contextually supplied set of actions.

1   Introduction

Our topic is ability modals, modals of the kind found in the following sentences:

(1)   John can go swimming this evening.

(2)   Mary cannot eat another bite of this rotten meal.

(3)   Louise is able to pick Roger up from work today.

As a simple heuristic, we can identify ability modals as modals that appear in sentences that can be paraphrased ⌜S is able to φ⌝ (on its most prominent reading) or ⌜S has the ability/power to φ⌝ (or with their negations).[1]

Our topic in particular is ability modals which have a specific action—an action tagged with a specific time—as the modal's prejacent, as in (1), (2), and (3). Other ability modals, as in

(4)   Susie can swim.

(5)   Jim is able to touch his nose with the tip of his tongue.

have as their prejacent a *generic* action, one not tied to a specific time. We assume that ability ascriptions with generic actions are just specific ability ascriptions embedded under a generic operator, and thus that a semantics for generic ability ascriptions will fall out of our proposal together with a suitable semantics for the generic operator.

We will survey three extant accounts of ability modals—the orthodox account, the universal account, and the conditional account—and argue that none is satisfactory. We then propose a new account that builds on the insights of the conditional account but avoids its well-known problems.

---

[1]Unless otherwise noted, all uses of 'can' and 'is able to' in examples and definitions that follow are to be read so that they can be paraphrased this way, and are to be treated as interchangeable.

2 The Orthodox Account

Kratzer (1977, 1981) gives a unified account of natural language modals according to which 'can' is an existential quantifier whose domain is the set of worlds that are 'best' according to a contextually supplied *modal base h* and *ordering source g*, both functions from worlds to sets of propositions. Given a world w, these together determine a preorder $\preceq_{g,w}$ on $\bigcap h(w)$ as follows: $u \preceq_{g,w} v$ iff $\{\psi \in g(w) : v \in \psi\} \subseteq \{\psi \in g(w) : u \in \psi\}$. Letting $\text{BEST}_{h,g,w}$ be the set of worlds in $\bigcap h(w)$ that are minimal according to $\preceq_{g,w}$, we have:

(6) **Orthodox Account**: $[\![\text{S can } \varphi]\!]^{h,g,w} = 1$ iff $\exists w' \in \text{BEST}_{h,g,w} : [\![\varphi]\!]^{h,g,w'} = 1$.

This account of modals is widely enough accepted that it has fair claim to being orthodoxy. But it's hard to see how to implement it for ability modals. The problem is that there is no natural value for the ordering source to take that makes the right predictions about ability modals.

On the standard implementation, the modal base is circumstantial, and the ordering source takes each world to a set of propositions that 'holds fixed certain intrinsic features of the agent in question' at that world.[2] But this approach makes predictions that are much too weak. Suppose Jim and Jo are at a crucial stage in a game of darts. Jo's young child Susie exclaims

(7)   Let me play! I can hit the bullseye on this throw.

But Susie can hardly ever hit the dartboard, and she has never even gotten a dart to stick in the dartboard; she can't, isn't able to, doesn't have the ability to, hit a bullseye on this throw. (To fix intuitions, imagine that Susie does go for the shot, and that the dart falls far short of the dartboard.) In light of these facts about Susie, (7) is clearly false on its abilitative reading. But on the approach just sketched, (7) is predicted to be true, since it is *compatible* with Susie's intrinsic properties, and local circumstances, that she hit a bullseye. (Note that (7) *does* have a true reading which can be paraphrased as

(8)   It can happen that I hit the bullseye on this throw.

This is a circumstantial/metaphysical reading of 'can'. The standard proposal thus adequately accounts for this reading, but not for the prominent abilitative reading of (7), paraphrasable not as (8) but as

(9)   I'm able to hit the bullseye on this throw.)

A natural first reaction to this issue is to include in the value of the ordering source at w propositions which describe what is *normal* at w. It's not obvious this helps even in this case: hitting the bullseye is unlikely but not obviously *abnormal*. But even granting that it is abnormal in a relevant sense for Susie to hit the bullseye, this proposal makes the wrong predictions in other cases: someone can be able to do something even if doing it is highly abnormal. For example, Susie is a competent speaker of English, and thus is able to utter the sentence

(10)   Every art, inquiry, action, and pursuit is thought to aim at some good.

But, being a small child and non-philosopher, she only utters this sentence in circumstances that are, intuitively, at least as abnormal as ones in which she hits the bullseye. The present

---

[2]Vetter (2013, p. 7); see Portner (2009). The following criticism also applies if *abilities* are in the ordering source.

proposal would thus wrongly predict that (11) is false:[3]

(11)   Susie can utter (10) now.

Incorporating normality into the ordering source thus will help only if we can spell out a notion of normality that treats Susie's hitting a bullseye as abnormal, but Susie's uttering (10) as normal. We don't see a natural way to do this, and thus we don't see a natural way for the orthodox semantics to account for ability modals (which isn't to say there is *no* way for the orthodox approach to account for ability modals; see Section 6).[4]

3   The Universal Account and the Dual

In response to worries like these, some have argued that abilitative 'can' has universal, rather than existential, force.[5] We do not think such an account could work.

   To show this, we need to take a brief excursus to discuss the dual of 'can'. Some have claimed that abilitative 'can' has no dual.[6] We find this puzzling. As with any other modal operator, we can form the dual of 'can' simply by putting a negation above and below it. Intuitions suggest that both 'cannot but' and 'cannot *not*' (italics indicating stress), as in

(12)   Ginger cannot but eat another cookie right now.

are natural language realizations of this semantic pattern, and thus are both duals of 'can'.[7] 'Must' and 'have to' can also have the meaning of the dual of 'can', as in

(13)   I have to sneeze. (Kratzer, 1977)

   Why have these data generally been ignored in accounts of ability modals? We believe an infelicity of nomenclature has promoted confusion here. 'Cannot but' is not an *ability* modal; rather, it expresses compulsion of some kind, whereas ability has to do not with compulsion but with potentiality. But this does not show that 'cannot but' is not the dual of 'can': it simply illustrates the humdrum point that duals do not have the same meaning as each other, even though they may have the same subject matter. Compare the situation with deontic 'may', which is a *permission* modal. Its dual does not express permission but rather (deontic) compulsion. But it does not follow that deontic 'may' lacks a dual; rather, 'may' and 'must' are simply part of a larger unified class, namely that of deontic modals.

   Likewise, 'can' and its dual do not both express something about ability, but they nonetheless belong to a larger unified class, which we propose to call the class of *practical modals*.[8] Carefully delimiting this class is important for the study both of natural language and of traditional philosophical problems. For instance, we suspect that the modals that appear in anankastic constructions are best analyzed as practical modals, as are the modals often adverted to in

---

[3]Portner (2009)'s suggestion to make 'can' a *good possibility* modal runs into the same problem.

[4]Kenny (1976) gives a different critique of the orthodoxy, pointing out that on a modal analysis, ⌜S can (φ or ψ)⌝ ⊨ ⌜S can φ or S can ψ⌝. The account we'll give inherits this objection. However, we are not persuaded by it; we believe that a supervaluationist approach along the lines of Stalnaker (1981) can answer this criticism.

[5]E.g. Kenny (1976, p. 139), Brown (1988), Giannakidou (2001), Giannakidou and Staraki (2012).

[6]See Hackl (1998). Hackl correctly notes that it is harder than we might expect to hear an equivalence between 'can' and, e.g. 'not must not'. We are not sure what to make of this fact, but we don't take it to show that ability modals 'lack a universal dual' (Hackl, 1998, p. 10).

[7]Thanks to Stephen Yablo for pointing out the first of these.

[8]Following a suggestion of Kieran Setiya (p.c.).

philosophical discussions of freedom of will and practical necessity.[9]

Once we have the dual of 'can' clearly in sight, it is hard to see how a universal analysis of 'can' could suffice. On such an analysis 'cannot but' would have existential force;[10] this cannot be squared with the intuition that 'cannot but' expresses compulsion of some sort.

4   The Conditional Analysis

A more promising approach than the orthodox or universal accounts treats ⌜S can φ⌝ as meaning the same as ⌜S would φ if S tried to φ⌝.[11] Let $f_c(w, \psi)$ be the selection function from Stalnaker (1968)'s theory of conditionals: a contextually supplied function from a world w and proposition ψ to the 'closest' world where ψ is true, and to w iff $w \in \psi$.[12] Let '$try(S, A)$' abbreviate ⟦S tries to do A⟧^{c}. Then:

(14)   **Conditional Analysis**: ⟦S can φ⟧^{c,w} = 1 iff ⟦$\varphi(S)$⟧$^{c, f_c(w, try(S,\varphi))}$ = 1.

φ ranges over actions, which we will model as properties; ⟦$\varphi(S)$⟧^{c} is the proposition that φ holds of S, i.e., that S does φ.

The conditional analysis (henceforth 'CA') gives much more intuitive truth-conditions than the orthodox account, and avoids the problems sketched for the orthodox account: e.g., if Susie tried to hit the bullseye, she wouldn't; but if she tried to utter (10), she would. But the CA has a number of problems of its own.

4.1   Monotonicity

First, the CA makes the wrong predictions about the monotonicity of 'can'. Intuitions about entailment suggest that 'can' is upward monotone:[13]

(15)   John can swim the butterfly tonight.

seems to entail

(16)   John can swim tonight.

Facts about NPI-licensing support this impression. 'Cannot' licenses NPIs, as in

(17)   John can't find any dance partners at this party.

According to a leading theory of NPI-licensing, the Fauconnier-Ladusaw analysis,[14] NPIs are licensed only in downward monotone environments; and 'not can' is downward monotone iff 'can' is upward monotone.

The CA, however, predicts that 'can' is non-monotone, not upward monotone. Some will take this to be a virtue; indeed, it's not clear whether a sentence like

(18)   I can ride a bike with training wheels right now.

---

[9]See e.g. von Fintel and Iatridou (2005), Aristotle (2002), Frankfurt (1969), Williams (1982), Setiya (2007).

[10]Surprisingly, Brown (1988) takes this prediction on board, claiming that the dual of 'can' is 'might'.

[11]An old philosophical idea, traceable to Hume (1748), taken up by Moore (1912) a.o.; for formulations in a model-theoretic framework, see especially Lehrer (1976), Cross (1986), and Thomason (2005). We use 'try' in an inclusive sense which treats things like cars and elevators as the kinds of things that can try.

[12]Stalnaker's uniqueness assumption is crucial for accounting for the negation and dual of 'can'.

[13]M is upward monotone iff if $\rho \vDash \psi$, then $M(\rho) \vDash M(\psi)$; downward monotone iff if $\rho \vDash \psi$, then $M(\psi) \vDash M(\rho)$.

[14]See e.g. Ladusaw (1979), von Fintel (1999).

entails

(19)   I can ride a bike right now.

We suspect that data like this can be explained pragmatically, however. For note that someone who can only ride a bike with training wheels could well insist

(20)   I can ride a bike with training wheels now, so *technically*, I can ride a bike now.

This discussion is inadequate to the complexity of the issues at hand, but we think the data suggest, *pace* the CA, that 'can' is upward monotone.

4.2   Dual

Second, the CA makes the wrong predictions about the dual of 'can':

(21)   **Dual (Conditional Analysis)**: $[\![S \text{ cannot but } \varphi]\!]^{c,w} = 1$ iff $[\![\varphi(S)]\!]^{c,f_c(w,try(S,\bar{\varphi}))} = 1$.

i.e. iff the closest world where S tries to not φ is one where S still φ's. This is too weak. Intuitively ⌜S cannot but φ⌝ means not only that S φ's if she tries not to, but that she φ's *no matter what she tries to do*. Another way to put the point is that the CA wrongly predicts that ⌜S cannot but *not* φ⌝ and ⌜S cannot but φ⌝ are consistent.

4.3   Counterexamples

Finally, the CA faces a number of related counterexamples.

First, there are cases in which the CA predicts that ⌜S can φ⌝ is true when it's false.[15] Suppose that John is planning to go to a movie alone. He has no special commitment to go—he has simply decided to. Ann asks John out to dinner; he replies:

(22)   I'm sorry, I'm not able to go; I'm going to a movie.

There is a prominent reading on which (22) is true. But if John *tried* to have dinner with Ann, he would succeed. So the CA cannot predict the true reading of (22). Importantly, 'able to' is abilitative on this reading, not deontic or bouletic: the retraction data associated with (22) differ from the retraction data we would expect if it were a deontic or bouletic modal. When someone makes a deontic or bouletic claim, rejoining with a claim about abilities feels like a non-sequitur. But in the present case, pointing out that John really does have an ability to go to dinner feels like a natural and effective response (as in (30) below), not a non-sequitur.

Second, there are cases in which the CA predicts ⌜S can φ⌝ is false when it's true.[16] Suppose Jones is a skilled golfer with an easy shot to make. Matt says:

(23)   Jones is able make this shot right now.

Matt has intuitively said something true. Now suppose Jones takes the shot and misses. We still judge Matt to have said something true. Afterwards, we can truly say

(24)   Jones (was able to/had the ability to) make that shot at that time.[17]

---

[15]Chisholm (1964), Lehrer (1968), Thomason (2005). Cf. cases with a phobic or comatose agent.

[16]From Austin (1961).

[17]There is also a false reading of (24), brought out when 'was able to' has perfective aspect (see Bhatt (1999) a.o.), but all that matters for our purposes is that there is a true reading, brought out when 'was able to' has imperfective aspect. Thanks to Nilanjan Das and Raphaël Turcotte for data in Hindi, Bengali, and French.

Yet given how f_c is defined, the closest world where Jones tries to make the shot is the actual world. Since Jones misses in the actual world, the CA thus wrongly predicts that (23) is false. Another way to put this point is that

(25)   Jones is able to make this shot right now, though if he tries of course he might miss.

is felicitous in some cases (as in this one), but is predicted by the CA to be always infelicitous.[18]

Finally, there are cases in which an agent can do something, but not if she tries to do it:[19]

(26)   David can breathe normally for the next five minutes.

is true, but if David tried to breathe normally, he would end up breathing abnormally.

From a technical point of view, these cases are easy to respond to: simply change the selection function relevant for ability ascriptions so that it selects worlds in a way that matches our intuitions.[20] The problem with this response is that it uncouples the CA from the analysis of conditionals, and thus from our intuitions about conditionals and similarity in general.[21] Without an intuitive characterization of the altered selection function, the resulting theory is not particularly predictive or explanatory. We believe we can do better.

5   Our Proposal

Examples like those discussed in the last section have been taken by many to refute the CA.[22] But we think that the CA is on the right track. It rightly captures the *hypothetical* nature of abilities: whether you are able to perform a particular action depends in some way on what happens under relevant alternate circumstances. Our account of ability modals aims to preserve this insight, but avoid the problems discussed in the last section by adding a layer of quantification over a contextually supplied set of actions to the meaning of 'can'.

Let $\mathcal{A}_{c,S}$ be a set of actions that are—in a sense to be precisified—*practically available* to an agent S in a context c. With f_c the selection function as above, we propose:

(27)   **Our Proposal**: $[\![\text{S can } \varphi]\!]^{c,w} = 1$ iff $\exists A \in \mathcal{A}_{c,S} : [\![\varphi(S)]\!]^{c,f_c(w,try(S,A))} = 1$.

I.e. $\ulcorner$S can $\varphi\urcorner$ is true just in case there is some contextually salient action A such that the closest world where S tries to do A is a world where S does φ.[23]

At a first pass, we may assume that in many cases, if an ability ascription has the form $\ulcorner$S can $\varphi\urcorner$, then $\mathcal{A}_{c,S} = \{\varphi, \bar{\varphi}\}$. In those cases the predictions of our account come very close to those of the CA. Thus, e.g., if we make this assumption in evaluating

(28)   Louis can go for a swim in the MIT pool tonight.

---

[18] We can make the inverse point as well. Suppose that young Susie by chance hit a bullseye; there is a reading on which 'Susie was able to make the bullseye in this game' remains false (a reading easiest to bring out with imperfective morphology), contrary to the predictions of the CA.

[19] See Vranas (2010) for discussion.

[20] Thomason (2005) suggests a response along these lines in response to (22).

[21] It is a non-negotiable property of similarity that nothing is more similar to something than itself, a thesis the altered selection function would have to abandon.

[22] See Austin (1961), Lehrer (1968), van Inwagen (1983).

[23] Chisholm (1964) makes a similar suggestion. As far as we know his suggestion hasn't been taken up in the subsequent literature, perhaps because he himself sketches a fairly serious objection to it; our account, however, avoids that objection by restricting the set of actions we quantify over.

then the action *going for a swim in the MIT pool tonight* is included in $\mathcal{A}_{c,Louis}$; thus provided that Louis can go swimming tonight if he tries to do so, we predict that (28) is true.

$\mathcal{A}_{c,S}$ won't always be set this way, however. We argue that flexibility in how this parameter is set by context crucially allows us to avoid the problems discussed in the last section: to avoid the counterexamples and make plausible predictions about the dual and monotonicity of 'can'.

5.1   Improved Predictions

The most important benefit of our view is that it avoids the counterexamples discussed in Section 4.3. Consider first cases in which ⌜S can φ⌝ is intuitively false, even though it is true that S would φ if S tried. Recall John, who says

(29)   I'm not able to go [to dinner]; I'm going to a movie.

Unlike the CA, our account can predict that (29) has a prominent true reading. To be sure, if John *tried* to go to dinner, he'd succeed. But, on our proposal, this does not guarantee the truth of 'John can go to dinner' at a context c: the action *meeting Ann for dinner* (or something much like it) must also be in $\mathcal{A}_{c,John}$. That is, it must be treated as *practically available* to John in a relevant sense in the context of utterance. If this condition is not met, then there is no action in $\mathcal{A}_{c,John}$ such that trying to do it guarantees that John meets Ann for dinner, and so (29) is true.

And indeed, we argue that this condition is not met in this case. An action may be treated as unavailable in a context for a wide variety of reasons. In this case, the fact that the agent has decided against the action suffices to rule it out. In other cases, an action may be treated as unavailable because it takes place in the past relative to the time of evaluation; because it violates certain rules that the agent in question is bound by; or because the agent is psychologically or physiologically unable to even consider it (as if John were phobic or comatose). Different notions of availability are at work in different contexts. Changes in the set $\mathcal{A}_{c,S}$ reflect changes in the relevant notion of availability.

Circularity threatens here. Among other things, facts about S's *abilities* might influence the constituents of $\mathcal{A}_{c,S}$. If we were aiming to give a reductive analysis of ability, we would want to identify the constituents of $\mathcal{A}_{c,S}$ without reference to abilities. But, like most projects of semantic analysis, our present aim is elucidation rather than reductive analysis (though we don't want to suggest that such a reductive analysis is impossible in the present framework).

One way to test the plausibility of our account is to see whether insisting on the availability of the action *meeting Ann for dinner* can modulate intuitions in this case. Suppose Ann replies:

(30)   Of course you *can* meet me—just skip the movie and come to dinner!

It seems that Ann has said something true: John *can* meet her for dinner. We hypothesize that, in responding this way, Ann ensures that $\mathcal{A}_{c,John}$ include *meeting Ann for dinner*, and so (30) comes out true.

$\mathcal{A}_{c,S}$ also plays a crucial role in responding to cases in which ⌜S can φ⌝ is intuitively true, even though it is false that S would φ if S tried. Recall the golf case. We said that

(31)   Jones can make this shot.

is intuitively true. How can we predict this? Suppose Jones aimed to the left of the pin; had he aimed to the right, he would have made the shot. Let the action *aiming to the right* be in $\mathcal{A}_{c,Jones}$. Then we predict that 'Jones can make this shot' is true even though Jones actually misses. This looks intuitive: indeed, we say of Jones

(32)   Well, he *could* have made the shot, if he had only aimed to the right.

This move can be applied quite broadly. We often ascribe abilities to agents even when they are not certain to succeed at a given action should they try, and even in cases where they fail when they in fact try.[24]

A worry about over-generation arises. Bob, a lousy chess player, is playing white in a chess match against Kasparov. Let $\mathcal{A}_{c,Bob}$ include, for all possible sequences of moves for white in the game, the action of taking that sequence of moves. One of these is such that, given Kasparov's course of play, if Bob completes that sequence of moves, he will beat Kasparov. It would seem to follow that

(33)   Bob can beat Kasparov in this match.

is true. But (33) is false (on its prominent reading).[25] To predict this, we propose to place certain default constraints on what goes into $\mathcal{A}_{c,S}$—roughly, that an action is in $\mathcal{A}_{c,S}$ only if the agent in question has the right kind of *epistemic* access to that action: perhaps, only if she knows that that action is a way for her to do φ. There is some sequence of moves that would beat Kasparov. But Bob doesn't *know which* sequence it is. This means that it is not in $\mathcal{A}_{c,Bob}$, which explains why we judge (33) false.[26]

Finally, appeal to $\mathcal{A}_{c,S}$ lets us explain why 'David can breathe normally for the next five minutes', and sentences like it, are true: there is *something* relevant (let us suppose) such that if David tries to do *that*, he breathes normally (say, *working on a problem set*).

The notion of practical availability that determines $\mathcal{A}_{c,S}$ is determined by context. As with other quantificational structures in natural language, change the context, and you change the domain of quantification, and thus the relevant notion of availability. Much more needs to be said to flesh out a complete meta-semantic story about how $\mathcal{A}_{c,S}$ is determined, and the plausibility of our view will turn on the ultimate success of such a project. But we take the present discussion to show that whether or not ⌜S is able to φ⌝ is taken to be true in a context c depends not just on what would happen if S tried to φ, but also on which ways, if any, of trying to φ are treated as practically available to S in c.

5.2   The Dual

Our approach makes good predictions about the meaning of the dual of 'can':

(34)   **Dual (Our Proposal)**: $[\![\text{S cannot but } \varphi]\!]^{c,w} = 1$ iff $\forall A \in \mathcal{A}_{c,S} : [\![\varphi(S)]\!]^{c,f_c(w,try(S,A))} = 1$.

Informally: for every action A available to S in c, S does φ in the closest world in which S tries to do A. In other words, no matter what S tries to do (among the actions we are treating as practically available in c), S ends up doing φ. Consider:

(35)   Ginger cannot but eat another cookie right now.

(35) says that Ginger is *compelled* to eat another cookie: *no matter what she tries to do*, she'll end up eating another one. This is precisely what we predict. And unlike the CA, our account of the dual rightly predicts that ⌜S cannot but *not* φ⌝ and ⌜S cannot but φ⌝ are inconsistent, provided that $\mathcal{A}_{c,S}$ is non-empty, a condition that, plausibly, is satisfied in most cases.

---

[24]One might be tempted to adopt a *graded* proposal to accommodate this kind of case. We do not see a fruitful way to spell out a view like that, though, and the present response seems satisfactory to us.

[25]Though making salient the sequence of moves in question can make prominent a true reading, as we predict.

[26]This also explains why flukey successes don't always merit ability ascriptions; see Footnote 18.

5.3 Monotonicity

Finally, 'can' is upward monotone on our account, matching intuitions and NPI-licensing.

6 Comparison and Conclusion

Our account is less of a departure from the orthodox Kratzerian approach to natural language modality than the CA. In keeping with orthodoxy, we say that 'can' denotes an existential quantifier. We depart from orthodoxy, however, in saying that 'can' quantifies over a set of actions (i.e. properties), not a set of worlds. That said, a number of recent proposals have departed from orthodoxy in a similar way, and converged on structurally similar accounts of natural language modals. Our account of ability modals differs from these because of the role that the selection function plays in the view. But it shares with them a central structural feature: quantification over sets of properties.[27]

It is too early to decide whether this approach will supplant orthodoxy, however, so it is also worth asking whether our view can be reframed within the orthodox framework. The answer is 'yes'. For any agent S, just let the modal base take any world w to $\varnothing$ and let the ordering source take any world w to $\{\{w'\} : \exists A \in \mathcal{A}_{c,S}(w' = f_c(w, try(S, A)))\}$. This account makes the same predictions as ours.[28] Should we adopt it?

We think not. If there were conclusive evidence in favor of the orthodox account for all other modals, then it would make sense to formulate our view in its terms. But short of that, we think there is no reason to go that way, because our approach fits far more naturally with a meta-semantic story.[29] A meta-semantics for both our view and its recasting in the orthodox framework must account for how $\mathcal{A}_{c,S}$ and f_c are determined, since both views refer to these parameters. But the orthodox recasting of our account *also* makes reference to (what looks to us like) a highly gerrymandered ordering source and an empty modal base. It is hard to see how this ordering source could play any role in an explanatory meta-semantics. So we think there is reason simply to dispense with it, as our account does.

We should also ask whether there are other elements of semantic machinery that can 'see' the differences between the semantic entries and thus distinguish them on more straightforward predictive grounds. We think there are; in particular, research into epistemic modals has suggested that embedded epistemic modals quantify over their *local contexts*.[30] Under this assumption, the behavior of epistemic modals embedded under practical modals may help decide between the two views under consideration.

The plausibility of our semantics will ultimately turn on this and a number of other open questions. In addition to those highlighted above, these include the question of whether we need to encode a non-accidental connection between S and φ in order for ⌜S can φ⌝ to come out true; how to account for graded ability ascriptions like 'Jones can easily make this shot'; and how to account for actuality implications that arise from the interaction of ability modals and aspect.

---

[27]See Yalcin (2012) and Willer (2013) on attitude verbs; Cariani et al. (2013) and Cariani (2013) on deontic modals; and Villalta (2008) on verbs of desire. Officially the quantification in these proposals is over sets of propositions, but this is a superficial modeling difference—propositions are just zero-place properties.

[28]*Modulo* embedded epistemic modals; see below.

[29]This is the same kind of choice point that will decide the parallel question in other cases, as between, for instance, the approach of Cariani et al. (2013) versus Kratzer's approach.

[30]Veltman (1996), Stephenson (2007), Yalcin (2007), Klinedinst and Rothschild (2012), Stalnaker (2014) a.o.

References

Aristotle (2002). *Nicomachean Ethics*. Oxford University Press, Oxford.

Austin, J. L. (1961). Ifs and cans. In *Philosophical Papers*, pages 151–180. Oxford University Press, London.

Bhatt, R. (1999). Ability modals and their actuality entailments. In Shahin, K., Blake, S., and Kim, E.-S., editors, *WCCFL 17 Proceedings*, volume 17, pages 74–87.

Brown, M. A. (1988). On the logic of ability. *Journal of Philosophical Logic*, 17:1–26.

Cariani, F. (2013). 'Ought' and resolution semantics. *Nous*, 47(3):534–558.

Cariani, F., Kaufmann, M., and Kaufmann, S. (2013). Deliberative modality under epistemic uncertainty. *Linguistics and Philosophy*, 36:225–259.

Chisholm, R. M. (1964). J. L. Austin's philosophical papers. *Mind*, 73(289):1–26.

Cross, C. B. (1986). 'Can' and the logic of ability. *Philosophical Studies*, 50(1):53–64.

von Fintel, K. (1999). NPI licensing, Strawson entailment, and context dependency. *Journal of Semantics*, 16(2):97–148.

von Fintel, K. and Iatridou, S. (2005). What to do if you want to go to Harlem: Anankastic conditionals and related matters. MS, MIT.

Frankfurt, H. G. (1969). Alternate possibilities and moral responsibility. *The Journal of Philosophy*, 66(23):829–839.

Giannakidou, A. (2001). The meaning of free choice. *Linguistics and Philosophy*, 24:659–735.

Giannakidou, A. and Staraki, E. (2012). Ability, action, and causation: from pure ability to force. In Mari, A., Beyssade, C., and Prete, F. D., editors, *Genericity*. Oxford University Press.

Hackl, M. (1998). On the semantics of 'ability attributions'. Master's thesis, MIT.

Hume, D. (1748). *An Enquiry Concerning Human Understanding*. Oxford University Press, Oxford.

Kenny, A. (1976). Human abilities and dynamic modalities. In Manninen and Tuomela, editors, *Essays on Explanation and Understanding*, pages 209–232. D. Reidel.

Klinedinst, N. and Rothschild, D. (2012). Connectives without truth tables. *Natural Language Semantics*, 20:137–175.

Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1(337-355).

Kratzer, A. (1981). The notional category of modality. In Eikmeyer, H. and Rieser, H., editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*. de Gruyter.

Ladusaw, W. (1979). *Polarity Sensitivity as Inherent Scope Relation*. PhD thesis, University of Texas at Austin.

Lehrer, K. (1968). Cans without ifs. *Analysis*, 29(1):29–32.

Lehrer, K. (1976). 'Can' in theory and practice: A possible worlds analysis. In Brand, M. and Walton, D., editors, *Action Theory*, pages 241–270. D. Reidel.

Moore, G. (1912). *Ethics*. Williams and Norgate, London.

Portner, P. (2009). *Modality*. Oxford University Press, New York.

Setiya, K. (2007). Cognitivism about instrumental reason. *Ethics*, 117:649–673.

Stalnaker, R. (1968). A theory of conditionals. In Rescher, N., editor, *Studies in Logical Theory*, pages 98–112. Oxford: Blackwell.

Stalnaker, R. (1981). A defense of conditional excluded middle. In Harper, W. L., Stalnaker, R., and Pearce, G., editors, *Ifs: Conditionals, Beliefs, Decision, Chance, and Time*, pages 87–105. D. Reidel.

Stalnaker, R. (2014). *Context*. Oxford University Press.

Stephenson, T. (2007). A parallel account of epistemic modals and predicates of personal taste. In Puig-Waldmueller, editor, *Proceedings of Sinn and Bedeutung*, volume 11, pages 583–97.

Thomason, R. H. (2005). Ability, action, and context. MS.

van Inwagen, P. (1983). *An Essay on Free Will*. Oxford University Press, New York.

Veltman, F. (1996). Defaults in update semantics. *Journal of Philosophical Logic*, 25(3):221–261.

Vetter, B. (2013). 'Can' without possible worlds: Semantics for anti-Humeans. *Philosophers' Imprint*, 13(16):1–27.

Villalta, E. (2008). Mood and gradability: an investigation of the subjunctive mood in Spanish. *Linguistics and Philosophy*, 31:467–522.

Vranas, P. (2010). What time travelers may be able to do. *Philosophical Studies*, 150:115–121.

Willer, M. (2013). Dynamics of epistemic modality. *Philosophical Review*, 122(1):45–92.

Williams, B. (1982). Practical necessity. In *Moral Luck*, pages 124–131. Cambridge University Press.

Yalcin, S. (2007). Epistemic modals. *Mind*, 116(464):983–1026.

Yalcin, S. (2012). Context probabilism. In *Proceedings of the 18th Amsterdam Colloquium*, pages 12–21.