Agentive Modals*

Matthew Mandelkern, Ginger Schultheis and David Boylan October 18, 2016

To appear, The Philosophical Review

Abstract

We propose a new analysis of a class of modals which we call *agentive modals*: ability modals and their duals, compulsion modals. After criticizing existing approaches — the existential quantificational analysis, the universal quantificational analysis, and the conditional analysis — we lay out a new account that builds on both the existential and conditional analyses. On our account, the *act conditional analysis* of agentive modality, a sentence like 'John can swim across the river' says that there is some *practically available action* (in a sense we make precise) which is such that, if John tries to do it, he swims across the river. We argue that the act conditional analysis avoids the problems faced by existing accounts of agentive modality, and show how it can be extended to an account of generic agentive modal claims. The upshot is a new vantage point on the role of agentive modal ascriptions in practical discourse: ability ascriptions serve as a kind of hypothetical guarantee, and compulsion ascriptions as a kind of non-hypothetical guarantee.

1 Introduction

We propose a new analysis of ability modals and their duals. Ability modals are modals like those used in (1), (2), and (3): modals that can be paraphrased with the dedicated ability modal 'able' or 'has the ability/capacity', on its most prominent reading.¹

- (1) John can go swimming this evening.
- (2) Mary can touch her nose with her tongue.
- (3) Louise is able to pick Roger up from work today.

^{*}We thank audiences at MIT, at the Semantics and Philosophy in Europe Eighth Colloquium at the University of Cambridge, and at the 20th Amsterdam Colloquium; anonymous reviewers and editors for this journal; three anonymous reviewers for the Amsterdam Colloquium, Alex Byrne, Fabrizio Cariani, Kai von Fintel, Martin Hackl, Samia Hesni, Sabine Iatridou, Abby Everett Jaques, Matthias Jenny, Justin Khoo, Joshua Knobe, John Maier, Eliot Michaelson, Milo Phillips-Brown, Paolo Santorio, Kieran Setiya, Bradford Skow, Robert Stalnaker, Zoltán Szabó, Timothy Williamson, and Stephen Yablo for valuable discussion and feedback.

¹Unless otherwise noted, we intend 'can' and 'is able to' – which we treat as interchangeable – to be interpreted agentively throughout what follows.

The modal constructions in (4), (5), and (6) are examples of the duals of ability modals, which we call compulsion modals:

- (4) Lara cannot but eat another cookie right now.
- (5) I have to sneeze right now.
- (6) I cannot *not* climb mountains.

We call the class of ability and compulsion modals taken together agentive modals.

We begin by discussing three existing accounts of agentive modals: the orthodox, existential account, according to which a sentence like 'John can go swimming this evening' means that *John goes swimming this evening* is compatible with a contextually determined set of worlds; the universal account, which holds that 'John can go swimming this evening' means that *John goes swimming this evening* is entailed by a contextually determined set of worlds; and the conditional account, which holds that 'John can go swimming this evening' means that if John tried to go swimming this evening, he would succeed.

We argue that none of these accounts is satisfactory. Along the way, we enumerate new desiderata for a satisfactory account of agentive modality. First, it must make the right predictions about the oft-neglected duals of ability modals, compulsion modals. This follows from the broader methodological point that a semantic theory for a certain expression must make plausible predictions not only about its meaning when it occurs unembedded, but also when it is embedded under operators like negation. In the course of this discussion, we identify two new classes of modals: the duals of ability modals (*compulsion modals*); and the class which contains both ability and compulsion modals (*agentive modals*). We suggest that these classes of modals play key roles in a broad range of philosophical debates.

Second, any satisfactory account of agentive modals must capture the *hypothetical* nature of ability ascriptions — ability ascriptions tell us what an agent *would* do under various alternative circumstances. But it must do so while avoiding a class of potent counterexamples to the conditional analysis.

Our account, the *act conditional analysis*, builds on the central insights of the orthodox analysis and the conditional analysis, while avoiding their problems. Our account treats ability ascriptions as a kind of *hypothetical guarantee*. When someone says 'John can go swimming this evening', she informs her interlocutors that *going swimming this evening* is, in a certain sense, within John's control. And we treat compulsion modals as a kind of *non-hypothetical guarantee*. When someone says 'I cannot but eat another cookie', she says that refraining from eating another cookie is *not* an option for her; she's guaranteed to eat another one. Our approach, furthermore, posits a central connection between agentive modals and two kinds of facts: *conditional* facts and *epistemic* facts. Whether an agent is said to be able to do (or refrain from doing) something depends on what would happen in some (possibly non-actual) scenario in which she tries

to do some relevant action. It also depends on whether we judge that the agent *knows how*, in some sense, to perform the action.

We spell these ideas out by arguing that a sentence like 'John can go swimming this evening' means that there is some *practically available* action (in a sense we make precise) such that if John tried to do that action, he would go swimming this evening. We thus capture the first of these intuitions, by positing that whether an agent can be said to be able to do an action depends on what would happen if she tried to do some (possibly different) action. We capture the latter intuition by positing that whether an action counts as practically available — and thus within the domain of quantification for agentive modals — partly depends on whether the agent knows that it is a way of doing the modal's prejacent.² Whether we treat an action as having this status depends, in part, on what kind of description of the agent's practical situation we have in mind: a more *objective* or more *subjective* description. This leads us to predict a distinction between *objective* and *subjective* readings of agentive modals — in parallel to the corresponding distinction between objective and subjective readings of deontic modals — a prediction which we argue is borne out.

In the first part of the paper we focus exclusively on *specific* agentive modal ascriptions: agentive modal ascriptions which have as their prejacent a *specific action* — an action indexed to a specific time — as in (1), (3), (4), and (5). We conclude the paper by enriching our account with a generic operator to make sense of *generic* agentive modal ascriptions — modal ascriptions like those in (2) and (6), which have as their prejacent a *generic* action, one not tied to a specific time. We show that our enriched account makes sense of the fact that many agentive modal claims say something general about what an agent is equipped to do (or refrain from doing) in normal circumstances.

2 The Orthodox Account

We start by summarizing and arguing against the orthodox account of agentive 'can', which traces back to Hilpinen (1969), is taken up in Lewis (1976), and spelled out in Kratzer (1977, 1981) (who we follow in our presentation). The orthodox account treats 'can' as 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 of these are functions from worlds to sets of propositions, functions which Kratzer calls *conversational backgrounds*. Given a world w, they together determine a set of 'best' worlds relative to h and g, BESTh,g,w. Then:³

(7) **Orthodox Account**:

²The *prejacent* of a modal is ambiguously the clause the modal embeds, or the denotation of that clause.

³Specifically, let $\preceq_{g,w}$ preorder $\bigcap h(w)$ as follows: $x \preceq_{g,w} y$ iff $\{\psi \in g(w) : y \in \psi\} \subseteq \{\psi \in g(w) : x \in \psi\}$. Then let BEST_{h,g,w} be the set of worlds in $\bigcap h(w)$ that are not strictly outranked by any other worlds, according to $\preceq_{g,w}$. For simplicity, we assume there is such a set (the *limit assumption*; see Lewis (1981), Stalnaker (1981) for discussion). Relaxing this assumption does not help with the difficulties we raise here. $\llbracket \varphi \rrbracket^{\langle p_1,\dots p_n,w \rangle}$ is φ 's extension relative to parameters $p_1 \dots p_n$ and world w; $\llbracket \varphi \rrbracket^{\langle p_1,p_2,\dots p_n \rangle}$ its intension.

$$[\![\mathbf{S}\ \mathsf{can}\ \varphi]\!]^{h,g,w}=1\ \mathsf{iff}\ \exists w'\in \mathtt{BEST}_{h,g,w}:[\![\varphi]\!]^{h,g,w'}=1.$$

Informally: $\lceil S \text{ can } \varphi \rceil$ is true just in case $\llbracket \varphi \rrbracket^c$ is true in some 'best' world; roughly, just in case $\llbracket \varphi \rrbracket^c$ is compatible with some contextually salient set of worlds.

This account is widely enough accepted as a general account of the meaning of natural language modals that it has fair claim to being orthodoxy. It has been successful in modeling a variety of 'flavors' of modality — deontic, epistemic, circumstantial — by varying the conversational backgrounds. But it is hard to see how to successfully implement this account in the case we are interested in: ability modals. Specifically, no natural value for the conversational backgrounds yields plausible predictions about the meaning of ability modals.

On the standard implementation of the orthodoxy for ability modals, 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 (Vetter (2013, p. 7); see also Portner (2009)). This approach predicts truth-conditions that are too weak.⁴ To see this, suppose Jim and Jo are at a crucial stage in a game of darts – Jo needs to hit a bullseye to win the game. Jo's overzealous young child Susie exclaims (8):

(8) Let me take your turn! I can hit the bullseye on this throw.

Susie hardly ever even *hits* the dartboard, and she has never gotten a dart to stick. Intuitively, then, (8) is not true in this scenario.⁵ We do not say that (8) is clearly false; its negation, 'Susie cannot hit the bullseye', does not strike us as clearly true either (a fact which our theory will try to make sense of by adverting to a notion of indeterminacy). What is clear is that (8) is not clearly true. However, on the approach just sketched, (8) is predicted to be clearly, determinately true, since it is certainly *compatible* with Susie's intrinsic properties, along with local circumstances, that she hit a bullseye.

Note that (8) *does* have a true reading which can be paraphrased as (9):

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

This is a circumstantial (or 'metaphysical') reading of 'can'. The present proposal adequately accounts for this reading, but not for the prominent agentive reading of (8), paraphrasable not as (9) but as (10).

(10) I'm able to hit the bullseye on this throw.⁶

⁴We will assume throughout what follows that the modal base is circumstantial, and consider changes to the ordering source; changing the modal base won't help with the problems we raise.

⁵To fix intuitions, imagine that Susie does throw the dart, and that it falls far short of the dartboard.

⁶A generalization seems to be lurking behind this test: in the semantic/syntactic sense, agentive modals are 'control' modals, whereas circumstantial modals are 'raising' modals. In the discussion of the conditional analysis and our own analysis, we will assume that the prejacent of agentive modals denotes an action, rather than a proposition; this assumption is made for concreteness, and does not have substantial import for either view. We do not mean to take a stand on the underlying question of the semantics of control (on which see e.g. Chierchia (1989), Brennan (1993), Wurmbrand (2002)).

On the most natural way of spelling it out, then, the orthodox account provides a suitable analysis of circumstantial modals, but not of ability modals.⁷

A natural first reaction to this case is to limit our domain of quantification to *normal worlds*, by including in the value of the ordering source for a world w a set of propositions describing what is *normal* at w. But this won't solve our problem. First, it's not clear that normality helps even in this case: although hitting the bullseye is unlikely, it is not obviously abnormal. But suppose we grant that it is abnormal, in a relevant sense, for Susie to hit the bullseye. The problem is that this proposal makes the wrong predictions in other cases, because agents are often able to do highly abnormal things. For example, Susie is a competent speaker of English, and thus is able to utter the sentence (11):

(11) The world is everything that is the case.

But, being a small child and non-philosopher, circumstances in which she utters this sentence are, intuitively, at least as abnormal as ones in which she hits the bullseye. The present proposal would thus wrongly predict that (12) is false:

(12) Susie can now utter 'The world is everything that is the case'.

Incorporating normality into the ordering source 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 (11) as normal. We don't see a natural way to walk this line.⁸

A different approach has abilities themselves determine what worlds count as best. On this approach, we include $[\![\varphi]\!]^c$ in the value of the ordering source at w iff $[\![S]\!]$ can $[\![\varphi]\!]$ is true at $[\!(c,w)\!]$. But this doesn't help with our problems: as long as hitting the bullseye is consistent with Susie's abilities, we will still predict that (8) is true. Suppose Susie has $[\![k]\!]$ (specific) abilities, of which at most $[\![n]\!]$ can be realized in any given world. Suppose further that for some $[\![n]\!]$ of her abilities, realizing all of them, together with local circumstances, is compatible with her hitting a bullseye. In that case, some best world will be one where Susie hits a bullseye, and so we will predict that (8) is true. Even if we let abilities directly determine the accessibility relation, then, we do not make the right predictions in the orthodox framework.

⁷At first glance, one might think the circumstantial modal (9) entails that Susie *can* hit the bullseye; after all, 'Susie might hit the bullseye, but Susie can't hit the bullseye' sounds marked. But this is not yet evidence that the inference from 'Susie might hit the bullseye' to 'Susie can hit the bullseye' is good: we will ultimately maintain that it is neither determinately true that Susie can or that she cannot hit the bullseye in this case, and likewise for cases like it, which independently explains the infelicity of conjunctions like this.

⁸Portner (2009)'s suggestion to make 'can' a *good possibility* modal (which diverges from the present proposal if we relax the limit assumption) faces the same objection: it cannot distinguish between cases like (8) and cases like (12). After all, *not uttering* (11) seems at least as 'normal' in the relevant sense as *uttering* (11), and so any world where Susie utters (11) will be outranked by a world where she doesn't. Even if this approach could avoid this issue, it makes the dual too weak: it predicts that $\lceil S \rceil$ cannot but $\varphi \rceil$ is true if e.g. there is an alternating descending infinite sequence of φ and $\neg \varphi$ worlds among the best worlds.

⁹Thanks to John Maier for pushing us on this point.

The orthodox semantics thus does not provide a fruitful framework for analyzing ability ascriptions. We suggest that this is because it aims to analyze them in terms of *compatibility*, which is an overly weak tool for analyzing ability.¹⁰

3 The Universal Account and Compulsion Modals

In response to worries like these, some have suggested that agentive 'can' has universal, rather than existential, force. We don't think such an account can work. To see why, we'll begin with a brief excursus on the dual of agentive 'can', which has been largely neglected in the literature. Indeed, some have claimed that agentive 'can' has no dual. A weak interpretation of this claim says that there is no *dedicated* lexical item that is uniquely interpreted as the dual of agentive 'can'. This seems true in English — all of the expressions which express the dual of agentive 'can' can also express other flavors of modality — but this fact is not particularly interesting for our purposes, nor is it unique to agentive 'can' (it seems to be true for many flavors of modality). A strong interpretation of the claim says that there is no way to express the semantic dual of agentive 'can'. This, however, is clearly false. As with any other modal operator, we can form the dual of 'can' by putting a negation above and below it; 'cannot but' and 'cannot *not*' (italics indicating stress), as in (13), realize this semantic pattern, and so are both duals of 'can': 13

- (13) Ginger cannot but eat another cookie right now.
- (14) Ginger cannot *not* eat another cookie right now.

'Must' and 'have to' can also have the meaning of the dual of 'can', as in (15) and (16), which seem to be equivalent to a paraphrase with 'cannot but':

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

(16) I must climb K2 this summer.

We believe that data like these have been neglected in part because of an infelicity of nomenclature. 'Cannot but' is not an *ability* modal. (13) does not say that Ginger is *able* to do something, but that she is *compelled* to do something. We propose to call modals with the meaning of 'cannot but' *compulsion modals*. This difference in meaning doesn't show that 'cannot but' is not the dual of 'can', however; it simply illustrates the point that duals do not have the same meaning as each other. Compare deontic 'may', a *permission*

¹⁰Which isn't to say there is *no* way for the orthodox approach to account for ability modals; see the conclusion. Kenny (1976) gives a different, influential critique of the orthodoxy; see Section 6.3.

¹¹E.g. Kenny (1976, p. 139), Brown (1988), Giannakidou (2001), Giannakidou and Staraki (2012).

¹²See Hackl (1998). The two versions of this claim have not, as far as we know, been carefully distinguished in the literature, which may have led in part to the neglect of the dual of 'can'.

¹³Thanks to Robert Stalnaker for pointing out the first of these.

modal. Its dual, deontic 'must', expresses *necessity*. 'May' and 'must' have different meanings but the same subject matter; together, they constitute a larger unified class: the class of deontic modals. Likewise, agentive 'can' and its dual have different meanings but the same subject matter: they too belong to a larger unified class, which we propose to call the class of *agentive modals*. ¹⁴ Carefully delimiting this class is important for the study of natural language and of traditional philosophical problems. For instance, we suspect that the strong necessity modals that appear in anankastic conditionals are best analyzed as compulsion modals; likewise for the modals adverted to in philosophical discussions of freedom and practical necessity. ¹⁵

Once we have the dual of 'can' clearly in view, we can quickly dispense with the universal analysis of 'can'. On that analysis, 'cannot but' will have existential force. This can't be squared with the intuition that 'cannot but' expresses compulsion. We do not see how we could derive this if 'cannot but' has existential force. The universal analysis might seem tenable when we only consider intuitions about unembedded ability ascriptions, but not when we turn our attention to ability ascriptions embedded under negation — illustrating the methodological importance of examining embeddings in semantic theory.

4 The Conditional Analysis

A more promising approach than either the orthodox or universal approach treats $\lceil S \rceil$ can $\varphi \rceil$ as meaning roughly $\lceil S \rceil$ would φ if $S \rceil$ tried to $\varphi \rceil.^{17}$ 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 $w \in \psi$. Let 'try(S,A)' denote the proposition $\llbracket S \rceil$ tries to do $A \rrbracket^c.^{18}$ Let φ range over linguistic items denoting actions, which we will model as properties, i.e., functions from individuals to propositions; let $\llbracket \varphi(S) \rrbracket^c$ be the proposition that $S \rceil$ does $\llbracket \varphi \rrbracket^c$.

Then:

(17) Conditional Analysis (CA):

¹⁴See e.g. Belnap (1991), Maier (2015b) for related but distinct usages.

¹⁵On the former see e.g. von Fintel and Iatridou (2005); on the latter, e.g. Aristotle (2002), Frankfurt (1969), Williams (1982), Setiya (2007).

¹⁶Brown (1988) takes this prediction on board, claiming, implausibly, that the dual of 'can' is 'might'. Brown's approach is formally similar to the conditional analysis, except that it does not make Stalnaker's uniqueness assumption. The present points tell against an approach like that, but a more sophisticated version is worth exploring, e.g. one which adopts a homogeneity constraint à la von Fintel (1997).

¹⁷This is 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).

¹⁸ Trying plays a key role in both the conditional analysis and our own analysis. We will not give a semantics for 'try' here (on which see e.g. Sharvit (2003) and citations therein). Because both the CA and our account analyze ability ascriptions in terms of trying and conditional reasoning, they make certain predictions about the relative rate of acquisition of certain concepts. For instance, these accounts predict that subjects will not be able to reason about abilities until they can reason about trying. This hypothesis is, as far as we know, compatible with the present state of empirical research, which suggests that children acquire the concept of trying, or purposive action in general, at a very early age; see e.g. Woodward (1998). Thanks to John Maier and Jonathan Phillips for discussion.

$$[S \operatorname{can} \varphi]^{c,w} = 1 \operatorname{iff} [\varphi(S)]^{c,f_c(w,\operatorname{try}(S,[\varphi]^c))} = 1.$$

On this approach, $\lceil S \text{ can } \varphi \rceil$ is true just in case S does $\llbracket \varphi \rrbracket^c$ in the closest world where S tries to do $\llbracket \varphi \rrbracket^c$; informally, just in case S would do $\llbracket \varphi \rrbracket^c$ if she tried to.¹⁹

The conditional analysis (henceforth 'CA') does better than the orthodox existential account. Recall Susie and the bullseye. Consider the conditional (18):

(18) If Susie tried to hit the bullseye now, she would.

Intuitively, (18) is not clearly true. That said, it does not strike us as clearly false either. We will follow Stalnaker (1981) in saying that the actual state of affairs, together with semantic facts, does not suffice to decide whether a conditional like this is true or false. A conditional like (18) is neither determinately true nor determinately false, but rather *indeterminate*.²⁰ Crucially, the CA predicts that (19) shares the truth value of (18).

(19) Susie can hit the bullseye now.

If (18) is not determinately true or determinately false, the CA predicts that (19) is not, either. This explains why we are both reluctant to accept (19) and its negation in (20):

(20) Susie can't hit the bullseye now.

The CA also rightly predicts that (12) is clearly true, since (22) is clearly true.

- (12) Susie can now utter 'The world is everything that is the case'.
- (21) If Susie now tried to utter 'The world is everything that is the case', she would.

Unlike the orthodox account, the CA is thus able to distinguish examples like (19) from examples like (12). It does this by capturing the fact that ability requires something more than compatibility: an ability is a kind of hypothetical necessity, a guarantee (or something close to it) that you will get something done if you try.

Despite its promise, the CA has a number of serious problems, which we turn to now.

4.1 Problems for Compulsion Modals

First, the CA makes the wrong predictions about *compulsion* modals:

¹⁹This last paraphrase depends on adopting Stalnaker (1975)'s theory of the conditional, which we remain non-committal on here (even though we draw heavily from it in giving our theory). If we do not adopt Stalnaker (1975)'s theory, then this paraphrase (and similar ones below) should be taken with a grain of salt.

²⁰A different approach, equally suitable for our purposes, treats (18) as *probably false* (Hawthorne, 2005).

(22) **Dual (Conditional Analysis)**:

$$[\![\mathbf{S} \text{ cannot but } \varphi]\!]^{c,w} = 1 \text{ iff } [\![\varphi(S)]\!]^{c,f_c(w,try(S,[\![\neg\varphi]\!]^c))} = 1.$$

That is, $\lceil S$ cannot but $\varphi \rceil$ is true iff the closest world where S tries to not $[\![\varphi]\!]^c$ is one where S does $[\![\varphi]\!]^c$; in other words, if S tries not to $[\![\varphi]\!]^c$, S does $[\![\varphi]\!]^c$.

This is too weak: $\lceil S \rceil$ cannot but $\varphi \rceil$ means not only that $S \rceil$ does $\llbracket \varphi \rrbracket^c$ if she tries not to, but that she does $\llbracket \varphi \rrbracket^c$ no matter what she tries to do. Another way to put the point is that the CA wrongly predicts that $\lceil S \rceil$ cannot but $\varphi \rceil$ and $\lceil S \rceil$ cannot but not $\varphi \rceil$ are consistent. For example, the CA predicts that 'Ginger cannot but eat another cookie' and 'Ginger cannot but not eat another cookie' are consistent. This strikes us as clearly false.

To bring this out, consider the following case.²¹ Suppose, unbeknownst to him, the buttons of the elevator in John's building have been incorrectly wired: if John presses the button marked '1' it takes him to the basement, and if he presses the button marked 'B' it takes him to the ground floor (suppose these are the only two buttons). John is in the elevator, and will press one of the two buttons. The CA predicts that (23) is true:

(23) John cannot but go to the first floor.

This is because, if John tries to not go to the first floor, then he will hit 'B', and thus end up on the first floor. For similar reasons, the CA predicts that (24) is also true:

(24) John cannot but go to the basement.

Both these predictions are wrong. (23) is false: after all, if John presses button '1', he will get to the basement. (24) is also false, for similar reasons. Furthermore, (23) and (24) are clearly inconsistent.

4.2 Problems for Ability Modals

Second, in a broad range of cases, intuitions about *ability* ascriptions come apart from intuitions about the corresponding conditional.

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

There is a prominent reading on which (25) is true — John can't make it to dinner because he's going to a movie. But (26) is clearly true:

²¹Thanks to Robert Stalnaker for suggesting a case like this to us.

(26) If John *tried* to have dinner with Ann, he would succeed.

The CA says that (25) is false if (26) is true, so it cannot predict this true reading of (25).

Importantly, there are a number of reasons to think that 'able' really is agentive in (25), and thus that the example cannot be dismissed on the grounds that the modal here is actually a deontic or bouletic modal. First, 'able' has a default agentive reading: out of the blue, speakers are inclined to hear 'able' as being about the agent's abilities, rather than about circumstantial, epistemic or deontic possibility. Second, we can easily set up the case so that John has no have any deontic commitment to go to the movie: simply suppose he is going by himself, and going simply because he wants to see the movie. Third, the retraction data associated with (25) are not what we would expect from a deontic or bouletic modal. When someone makes a deontic or bouletic claim, rejoining with a claim about ability 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 to (25) (as in (36) below).

Examples like (25) are widespread; speakers often use $\lceil I \operatorname{can}' t \varphi \rceil$ or $\lceil I' \operatorname{m}$ not able to $\varphi \rceil$ to communicate that they have a commitment which conflicts with doing $\llbracket \varphi \rrbracket^c$. A semantics for 'can' must capture this fact. The CA does not.

A structurally similar case comes from Lehrer (1968).²² Suppose that Larry is offered a bowl of red candy. He has a pathological phobia of red candy; nothing could induce him to take such a candy. (27) seems false.

(27) Larry is able to take a candy.

But, one might think, if he tried to take the candy, he'd succeed — the closest world where he tries to take the candy is one where he does not have the phobia in question. So, one might think, the CA wrongly predicts that (27) is true.

We are not completely convinced that (28) is true in this situation.

(28) If Larry tried to take the candy, he would.

After all, Larry has a phobia! (See Albritton (1985) for attempts to dismiss cases like this along these lines.) But we are happy to grant the judgment that (28) is true. If this is right, it's another problem for the CA which, as we'll show, our account does not have; if not, examples like (25) already provide decisive reason to reject the CA.

The examples we have discussed so far are counterexamples to the CA as a *sufficient* condition for the truth of an ability ascription. There are also counterexamples to the CA as a *necessary* condition for the truth

²²See also Chisholm (1964); a variant has a comatose agent Len. We want to say that 'Len can take a walk' is false; but if Len tried to take a walk, he would be conscious, and so he would in fact take a walk.

of ability ascriptions — cases in which $\lceil S \text{ can } \varphi \rceil$ is true but $\lceil S \text{ would } \varphi$ if she tried is false. Consider the following case, adapted from Austin (1961). Jones is a skilled golfer with an easy shot onto the green. Matt says:

(29) Jones is able to make this shot right now.

Matt has said something true. Now suppose Jones takes the shot and misses the green. We may still judge Matt to have said something true; afterwards, we can truly say (30):

(30) Jones (was able to/had the ability to) make that shot at that time.²³

Yet given how the selection function is defined, the closest world where Jones tries to make the shot is the actual world. Since Jones actually misses, the CA wrongly predicts that (29) is false. A way to press this point is to point out that sentences like (31) are often felicitous:

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

The CA predicts that the two conjuncts in (31) are incompatible, since $\lceil \text{If p}$, then $q \rceil$ and $\lceil \text{If p}$, then might not $q \rceil$ are incompatible.

A final kind of case that poses problems for the CA is one in which an agent can do something, but only if she does not try to do it, as in (32):²⁴

- (32) David can breathe normally for the next five minutes.
- (32) is true if David is healthy. But if David *tried* to breathe normally, he'd concentrate on his breathing and end up breathing abnormally.

From a technical point of view, all of the cases in this section are easy to respond to: simply choose a selection function for ability ascriptions that selects worlds in a way that matches our intuitions.²⁵ 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.²⁶ Without an intuitive characterization of the altered selection function, the resulting theory is not predictive or explanatory.

²³There is also a false reading of (30), 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, which is clearly brought out in languages that mark aspect morphologically, when 'was able to' has imperfective aspect. Thanks to Nilanjan Das and Raphaël Turcotte for data in Hindi, Bengali, and French.

²⁴See Vranas (2010) for discussion. Again, not everyone accepts these judgments. We find this case compelling, but our rejection of the CA does not depend on these judgments.

²⁵Thomason (2005) suggests a response along these lines.

²⁶It is a non-negotiable property of similarity that nothing is more similar to something than itself, a thesis we would have to abandon if we adopted the altered selection function to make sense of cases like (29).

5 The Act Conditional Analysis

Examples like those discussed in the last subsection have been taken by many to refute the CA.²⁷ 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. We give an account of agentive modals that aims to preserve this insight, while avoiding the problems faced by the CA.

Our account accomplishes this by incorporating into the meaning of 'can' a layer of quantification over a contextually supplied set of actions. Let $\mathcal{A}_{S,c}$ be a set of actions which are - in a sense to be precisified - practically available to an agent S in a context c. Where f_c is again Stalnaker's selection function, and φ denotes, again, a specific action:

(33) Act Conditional Analysis (ACA):

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

According to the *act conditional analysis* (henceforth 'ACA'), $\lceil S \operatorname{can} \varphi \rceil$ is true just in case there is some practically available action A such that the closest world where S tries to do A is a world where S does $\llbracket \varphi \rrbracket^c$ (note that A needn't be $\llbracket \varphi \rrbracket^c$). In other words, there is some practically available action such that if S tries to do it, she does $\llbracket \varphi \rrbracket^c$. 28,29

At a first pass, we may assume that in many cases, if an ability ascription has the form $\lceil S \text{ can } \varphi \rceil$, then $\llbracket \varphi \rrbracket^c \in \mathcal{A}_{S,c}$. In those cases the predictions of the ACA come very close to those of the CA. So, for example, if we make this assumption in evaluating (34), then the action *going for a swim in the pool tonight* is included in $\mathcal{A}_{Louis,c}$.

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

In this case, (34) is true if Louis will go swimming tonight should he try to do so. The similarity of the ACA to the CA allows it to inherit the main virtue of the CA sketched above: it is able to distinguish between cases like (19) and (12), repeated here:

- (19) Susie can hit the bullseye now.
- (12) Susie can now utter 'The world is everything that is the case'.

²⁷See, among others, Austin (1961), Lehrer (1968), van Inwagen (1983).

²⁸Chisholm (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; the ACA, however, avoids that objection by restricting the set of actions we quantify over. See also Albritton (1985), Setiya (2007), and Maier (2015b) for discussions with some connection to the present proposal.

²⁹We leave open whether we need to encode a non-accidental connection between S and $[\![\varphi]\!]^c$ in order for $[\![S]\!]$ S can $[\![\varphi]\!]$ to come out true. We are inclined to think that the connection is implicated rather than encoded.

Assuming for simplicity that $A_{Susie,c}$ includes just the actions throw a dart towards the bullseye and utter 'The world is everything that is the case', and their complements, we predict that (12) is clearly true, since if Susie tries to utter 'The world is everything that is the case', she succeeds. On the other hand, we predict that (8) is not clearly, determinately true, since none of these actions is determinately such that if Susie tries to do it, she succeeds.

As we will now show, in addition to capturing these attractive predictions of the CA (and thus avoiding our objections to the orthodox account), the ACA also avoids the problems for the CA raised above.

5.1 Compulsion Modals

First, our approach makes plausible predictions about the meaning of compulsion modals:

(35) **Dual (Act Conditional Analysis):**

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

Informally: for every action A practically available to S in c, S does $\llbracket \varphi \rrbracket^c$ 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 $\llbracket \varphi \rrbracket^c$. This looks like a plausible prediction — much more plausible than the CA's prediction that $\lceil S$ cannot but $\varphi \rceil$ is true just if S does $\llbracket \varphi \rrbracket^c$ provided S tries not to. Consider (13) again:

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

Intuitively, (13) says that Ginger is *compelled* to eat another cookie: *no matter what she tries to do*, she'll eat another one. This is precisely what we predict. Moreover, the ACA rightly predicts that $\lceil S \rceil$ cannot but $\varphi \rceil$ and $\lceil S \rceil$ cannot but *not* $\varphi \rceil$ are inconsistent:³⁰ if $S \rceil$ does $\llbracket \varphi \rrbracket^c$ no matter what she tries, then it's not the case that $S \rceil$ does $\llbracket \neg \varphi \rrbracket^c$ no matter what she tries.

Note that the plausibility of our predictions here stems in part from our choice to use the selection function from Stalnaker (1968), which selects a single world;³¹ had we followed Lewis (1973) in adopting a selection function which selects a *set* of worlds, we would make the implausibly weak prediction that $\lceil S \rceil$ cannot but $\varphi \rceil$ is true just in case, for every practically available action, if S tries to do that action, then she *might* do $\llbracket \varphi \rrbracket^c$.

5.2 Counterexamples Avoided

Second, crucially, the ACA is able to avoid all the counterexamples to the CA discussed in Section 4.2. Consider first cases in which $\lceil S \operatorname{can} \varphi \rceil$ is intuitively false, even though it is true that S would $\llbracket \varphi \rrbracket^c$ if S tried.

³⁰Provided that $A_{S,c}$ is non-empty, a condition we assume is met in most cases.

³¹Though note that this does not commit us one way or the other on the semantics of the conditional.

Recall John, who says (25):

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

Unlike the CA, the ACA can predict that (25) 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 that 'John can go to dinner' is true: the action *meeting Ann for dinner* (or something much like it) must also be treated as practically available to John. If this condition is not met, then there is no action in $\mathcal{A}_{John,c}$ such that trying to do it guarantees that John meets Ann for dinner. In that case, we predict that 'John can go to dinner' is false, and that (25) is true.

We conjecture that *meeting Ann for dinner* is not treated as practically available in this context because John has decided against this action. One way to test the plausibility of this hypothesis is to see whether insisting on the availability of the action *meeting Ann for dinner* can modulate intuitions: in general, speakers tend to defer to an insistence on a larger domain of quantification. Suppose Ann replies:

Of course you're *able to* meet me - just skip the movie and come to dinner!

It seems that Ann has said something true; after hearing (36), we are inclined to judge that it is true that John can meet Ann for dinner. By uttering (36), Ann ensures that $A_{John,c}$ include *meeting Ann for dinner*, and so (36) comes out true.

In a moment we will discuss, in more disciplined and general terms, what practical availability amounts to. For the present, though, we will show how flexibility in how this parameter is determined lets us respond in a similar way to the other cases discussed in Section 4.2.

Recall the case in which Larry is offered a bowl of candy, but his phobia prevents him from taking one. If Larry is so phobic that we cannot really entertain the possibility of his trying to take the candy, then we may well not treat *taking the candy* as practically available for Larry. In that case, we will predict that 'Larry can take the candy' is false, even if 'If Larry tried to take the candy, he would' is true. Likewise if Len is comatose, then we may not treat, e.g., *going for a walk* as practically available for him.

We can make similar moves in response to cases where $\lceil S \operatorname{can} \varphi \rceil$ is true, even though it is false that S would φ if S tried. Recall the golf case. We said that (29) has a true reading.

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

Now suppose Jones aimed to the left of a certain tree; because of a freak gust of wind, the ball was blown off course. Had he tried to aim to the right, he would have made the shot. Let *aiming to the right* be in $\mathcal{A}_{Jones,c}$. Then we predict that (36) is true even though Jones actually misses. This nicely captures intuitions: it is natural to say (37) about Jones.

(37) Well, he *was* able to make the shot; all he had to do was aim to the right.

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. The ACA nicely accounts for this practice. (We discuss cases like this further in Section 6.2 below.)

Finally, appeal to $A_{S,c}$ lets the ACA explain why (38) and its ilk strike us as true.

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

There is *something* relevant (we may suppose) such that if David tries to do *that*, he breathes normally (say, *working on a paper*), and thus the ACA predicts that (37) is true.

5.3 Practical Availability

We have shown that, by incorporating quantification over a set of actions, the ACA can avoid our counterexamples to the CA, without uncoupling the analysis of agentive modals from intuitions about conditional facts. But unless we say more about how this set of actions is determined — i.e., what practical availability amounts to — we face a charge of *ad hoc* maneuvering. In this section we give a general characterization of practical availability that regiments the intuitions elicited in the last section. We do not aim to provide a universal characterization of practical availability; we believe it is a genuinely context-dependent notion, and as such may be determined in a variety of different ways in different contexts. But we hope to say enough to give a sense of speakers' default way of thinking about practical availability, and thus to make the predictions of our account more concrete.

Consider the following case, due to Timothy Williamson (p.c.), which will help bring out the constraints we need to place on practical availability. Imagine a grid with 100 buttons, labelled 1-100. Exactly one of these buttons is such that, if Lizzie pushes it, she wins a prize, but she doesn't know which. She gets one shot at winning the prize. It seems clear that (38) has a prominent agentive reading on which it's not (determinately) true; after all, Lizzie doesn't know which button to push in order to win the prize:

(38) Lizzie is able to win the prize.

This judgment is important to get clearly in mind; one way to bring it out is to imagine Lizzie in a high-stakes situation, say in a rapidly descending plane. One of 100 buttons on the console turns on autopilot; the others ignite the fuel. The panicked crew-members ask:

(39) Can anyone here turn on autopilot?

If Lizzie replied with (40), the crew-members would rightly hold her to account if they learned that she doesn't know which button enables autopilot.

(40) I can!

On their most prominent readings, neither (40) nor (38) is determinately true (we do not hold that they are determinately false either; they may be indeterminate).

It's clear what actions should count as practically available in order for us to predict these judgments. Consider again (38). We want to say that $A_{Lizzie,c}$ contains the actions winning the prize, not winning the prize, but not the more specific actions pushing button one, pushing button two, and so on. If $A_{Lizzie,c}$ did contain these more specific actions, then we'd predict that (38) is true. Suppose the seventh button is the winning one; it's clearly true of Lizzie that if she tried to push the seventh button, she would succeed. On the other hand, if $A_{Lizzie,c}$ contains just the actions winning the prize, not winning the prize, then we predict that (38) is not determinately true, since there will be no practically available action which is determinately such that if Lizzie tried to do it, she would win the prize.

But how do we predict that this is how $A_{Lizzie,c}$ is set? This question looks particularly pressing when we compare our judgments about (38) to our judgments about (41).

(41) Lizzie can press the k^{th} button.

It seems clear that every instance of (41) is true, for $1 \le k \le 100$. In order for these to be true, however, something like *pressing the* k^{th} *button* must be treated as practically available, for every k. But then we will predict that (38) is true, since there will be some practically available action such that if Lizzie tries to do that, she wins.

We can make the right predictions about both (38) and (41) by treating $A_{S,c}$ as sensitive to the prejacent of the modal (which is, after all, part of the contextually available information). In particular, we suggest that an action ψ will typically count as practically available for S in a context c just in case S could reasonably conclude in favor of doing ψ with respect to the goal of doing the prejacent or its complement.³² This is not meant as a strict rule; actions might count as practically available even when they are not reasonable in this sense (for instance, after the fact, the action that an agent actually does will almost invariably count as practically available, whether or not it was reasonable). But we believe reasonableness provides a helpful general heuristic for fleshing out the set of practically available actions.

Reasonableness will involve many things. Among them will be a certain kind of epistemic standing. We propose that an action A is reasonable in the relevant sense only if the agent knows that A is a way to bring about the prejacent — or knows that it is a way to bring about its complement — relative to a certain description of her practical situation (more on this below).³³ In other words, the agent must know that, given that description, if she does A, then she does the prejacent; or that if she does A, she does the prejacent's complement.

³²We include the complement as a goal in order to make correct predictions about negated ability modals.

³³Perhaps (true) belief would suffice; we do not have strong committments on this point.

Applying this idea to our case: relative to the goals of winning the prize, not winning the prize – the prejacent of (38) and its complement – and assuming that we take Lizzie's knowledge of her practical situation to provide the relevant description of her practical situation, the set of practically available actions will just be winning the prize (plus any actions that Lizzie knows are ways of not winning the prize). In particular, pressing the seventh button will not be treated as practically available, since Lizzie does not know that this is a way to win the prize (or to not win the prize). By contrast, she of course does know that winning the prize is a way of winning the prize, so, assuming that this action is otherwise a reasonable thing for Lizzie to conclude in favor of, it will be treated as practically available. Thus (38) will not be determinately true, since no practically available action is determinately such that if Lizzie tries to do it, she wins the prize. By contrast, relative to the action pushing the k^{th} button (the goal made salient by the prejacent of (41)), pushing the k^{th} button will itself be treated as practically available, since Lizzie knows this is a reasonable way of pushing the k^{th} button. So we predict that (41) will be true in its context.

Our approach thus lets us predict a true reading of (41) and an indeterminate reading of (38). So far so good. But in addition to an indeterminate reading, (38) also seems to have a true reading. This is easiest to bring out if we consider a bystander commenting on the situation; the bystander could say (42):

(42) Lizzie can win the prize; she just has to push the seventh button.

How can we predict this? We have said that an action counts as practically available only if the agent knows that it is a way of bringing about the prejacent *relative to a given description of her practical situation*. There are different ways of supplying the relevant description. On the *subjective* way of thinking about it we have assumed just now, the description of Lizzie's practical situation is limited to what she *knows* about her actual situation. We think this reading is the default reading; it can be brought out by focusing on what makes sense given an agent's limited information, e.g. by making salient how Lizzie thinks about her situation:

(43) Lizzie can't put the plane on autopilot – she has no idea which button to press!

But on a more *objective* way of thinking about it, the description of Lizzie's practical situation is just a complete description of all the facts relevant to her *actual* situation. An objective reading is elicited by making salient the actual facts at hand, as in (42).

The epistemic constraint on practical availability can thus be read either objectively or subjectively. Somewhat surprisingly, then, agentive modals, like deontic modals, have both objective and subjective readings.³⁴

The present characterization of practical availability helps answer a possible worry about our view. Suppose again that the seventh button is the winning button. On our view, 'can' is upward monotone: if

³⁴Our structural approach to the distinction parallels Cariani et al. (2013)'s approach to deontic modals. We assume this distinction admits of gradations tracking different ways of specifying a practical situation.

 φ entails ψ , then Γ S can $\varphi \Gamma$ entails Γ S can $\psi \Gamma$.³⁵ Since in this context *pushing the seventh button* entails *winning the prize*, we predict that (44) entails (38) (repeated here). Thus the falsity of (38) should entail that (44) is false, too.

- (44) Lizzie can push the seventh button on the grid.
- (38) Lizzie is able to win the prize.

But (44) sounds clearly true. This is a *prima facie* puzzle for any upward monotone account, like ours, and indeed, as an editor for this journal points out, for any account which validates the substitution of contextually equivalent prejacents under agentive modals – since in this case, *pushing the seventh button* is contextually equivalent to *winning the prize*.³⁶

Our characterization of practical availability can resolve this puzzle, however. Though we do predict that (44) is sometimes false (whenever (38) is), it does not follow that we can ever *hear* a false reading of (44). On our account, false instances of (44) are *elusive* in the sense of Lewis (1996): any context in which (38) is not true is one in which (44) is not true, but the moment we *entertain* (44), the context changes in a way that makes (44) come out true, since by changing the prejacent, we change the set of practically available actions.

Moreover, it looks like we *need* upward monotonicity — or at least the substitutability of contextual equivalents — to explain certain reasoning patterns elicited here. Suppose Matt says that Lizzie is unable to win the prize. Out of the blue, we're inclined to agree with Matt. But now suppose that Sally responds as follows:

(45) Lizzie can push the first button, the second, and so on. But one of these is the winning button. So she *can* push the winning button, and thus she's able to win!

It would be reasonable for Matt to retract his claim after hearing Sally's case. Once we recognize that Lizzie can push the winning button, whichever it is, we feel compelled to admit that she win the prize. Our account captures this fact, while also making sense of the difference in intuitions about the truth of (38) versus (44).

We conclude this section by showing that our characterization of practical availability regiments and explains the assumptions we made in showing how the ACA avoids the counterexamples to the CA. Recall

³⁵This generally matches intuitions, *modulo* 'free choice effects' (Kamp, 1973) which must be accounted for; more thorough discussion of monotonicity must await future work.

³⁶To be precise, what we need to get the problem going is intersubstitutability of *contextual* equivalents, or upward *contextual* entailment, which both differ from upward *logical* entailment, which is all that really follows from our view. But, under the standard assumption that the relevant selection function always selects a world in the context set when there is one which verifies its propositional argument (Stalnaker, 1968), then both those principles will indeed follow from our view whenever, for every practically available action, it is compatible with the context that the agent in question tries to do that action — a condition which seems to be met in the present case, and indeed in many cases. A different reaction to this case would be to relax that standard assumption about the selection function.

John, who has decided to go to a movie. Is there a reasonable action for him to conclude in favor of doing, relative to the pair of goals *going to dinner with Ann*, *not going to dinner with Ann*, such that if he tries to do it, he goes to dinner with Ann? It is easy to get the intuition that there is not: since he has decided to go to a movie, it would not be rational for him to deliberatively conclude in favor of something inconsistent with that. People should stick with their decisions. On the other hand, one can also get the intuition (with slightly more difficulty) that there is a reasonable action for John to conclude in favor of doing, relative to the pair of goals *going to dinner with Ann*, *not going to dinner with Ann*, such that if he tries to do it, he goes to dinner with Ann. After all, John could simply cancel his movie plans; people should be willing to revisit decisions. Depending how we come down on this, we will treat *going to dinner with Ann* as practically available (in the latter case) or not (in the former), and thus predict both true and false readings of (46).

(46) John can go to dinner with Ann.

We can say similar things about phobic Larry. Taking candy would not be a reasonable course of action for phobic Larry to deliberatively conclude in favor of, for obvious reasons, and thus will not be treated as practically available. (One might worry about circularity in this account of practical availability. Among other things, facts about S's *abilities* might influence what counts as a reasonable course of deliberation; indeed, this seems to the case here. But, like most projects of semantic analysis, our aim is elucidation, not reductive analysis.)

This discussion brings out an important fact about our characterization of practical availability, worth emphasizing here: the epistemic constraint we discussed above is only meant to be a *necessary* condition on practical availability, not a sufficient one. In addition to the epistemic constraint, more general considerations about reasonableness will also influence whether an action is treated as practically available.

Recall Jones the golfer. Suppose that Jones did not know that aiming just to the right of the tree would result in a successful shot, because he did not know that a gust of wind would arise at just the moment he made the swing. Jones did know, however, that, given such a gust of wind, aiming just to the right of the tree would result in a successful shot. So, relative to an *objective* description of his practical situation, *aiming just to the right of the tree* is practically available. Assuming that, had he tried to aim just to the right of the tree, he would have made the shot, we thus predict a true, objective reading of (29).

(29) Jones is able to make this shot right now.

This objective reading can, again, be brought out by emphasizing the facts of Jones' situation - facts he might not be aware of, as in (37), repeated here:

(37) Well, he *was* able to make the shot; all he had to do was aim to the right.

By contrast, for a less skilled golfer, Milo, the analogue of (29) will not be true, even on an objective reading. For even if Milo knew all of the relevant facts about his practical situation, it is likely that, if Milo tried to aim to the right, he would *still* not make the shot. Thus the truth of (29) depends in an important way on the fact that Jones is a skilled golfer.

One might wonder whether the distinction between a subjective and objective reading of agentive modals which we have spelled out in this section, and which we are relying on in our analyses of these cases, is really necessary. An alternate hypothesis is that what we have been calling an objective reading of an agentive modal is really just a circumstantial modal.³⁷ We think, however, that there is an important distinction to be made here: while the objective agentive reading may well entail the circumstantial reading in most cases (since it makes salient a world where the prejacent happens), the converse is not true.

For instance, take a variant on the Susie case. Susie is now playing against Smith, an expert dart player who is extremely reliable in making bullseyes at this distance. What they don't know is that they are playing with darts that are slightly heavier than usual; to make a bullseye, they would have to compensate for this.³⁸ Jo, who is watching the game, says:

Smith is able to hit a bullseye now - he just has to compensate for the weight.

This sounds perfectly true here. This is what we have been calling an objective reading of 'able'; the subjective reading is false, since Smith doesn't know that he needs to compensate for the weight of the bullseye. By contrast, (48) does not sound (determinately) true; the situation here is identical to the original case involving Susie.

(48) Susie is able to get a bullseye now — she just has to compensate for the weight.

But paraphrases eliciting the circumstantial reading pattern differently:

- (49) It can happen that Susie hits the bullseye.
- (50) It can happen that Smith hits the bullseye.

Both (49) and (50) sound true. Thus it is circumstantially possible both for Smith and for Susie to hit a bullseye. But if the prominent true reading of (47) really is a circumstantial reading, as the present objection goes, then we should expect a prominent true reading for (48). This is not borne out. Nor can a subjective agentive reading predict this contrast between (47) and (48), since, again, (47) is false on a subjective agentive reading. By contrast, the objective reading of agentive modals which we have posited predicts precisely the contrast we observe here. Since Smith is highly skilled, there is something which (determinately) would

³⁷Thanks to an anonymous reviewer for this journal for pushing us on this question.

³⁸Note that Smith's predicament is similar to that of Jones the golfer.

have resulted in him making the shot, if he had tried it — he simply does not in this case know what action that is. Not so for Susie: since she is a small child, no action is (determinately) such that, had she tried it, it would have resulted in her hitting the bullseye. We thus must distinguish circumstantial readings from objective agentive readings, which, in turn, we distinguish from subjective agentive readings.

5.4 Interim Summary

This completes our exposition of the ACA. Putting the pieces together: our account says that $\lceil S$ can $\varphi \rceil$ is true iff there is some practically available action A such that if S tries to do A, S does $\llbracket \varphi \rrbracket^c$. Practical availability is, in turn, a context-dependent notion. We have proposed the following rough heuristic for thinking about it: we say that an action A is practically available to an agent S with respect to a salient goal ψ iff S could reasonably conclude in favor of doing A in order to do ψ . When it comes to ability ascriptions we assume that the salient goals are typically the prejacent of the ascription and its complement, and we assume, *inter alia*, that S can reasonably conclude in favor of doing A in order to do ψ only if she knows that - given a certain description of the world - if she does A, she does ψ . Our characterization of practical availability is somewhat vague. But so are ability ascriptions, and the vagueness of the two seem to go hand in hand.

Our account answers to an important intuition about what ability ascriptions are used to do: when someone says $\lceil I \operatorname{can} \varphi \rceil$, she is assuring her interlocutors that $\llbracket \varphi \rrbracket^c$ is within her control in a certain way; this is why we do not judge that Susie can hit the bullseye, or that Lizzie can turn on autopilot. This is precisely what the ACA predicts: $\lceil S \operatorname{can} \varphi \rceil$ says that there is some relevant action — some action that S could reasonably deliberate about and conclude in favor of doing — such that if S tries to do it, S succeeds in doing $\llbracket \varphi \rrbracket^c$.

Ability ascriptions, on our account, serve as a kind of hypothetical guarantee. This is a strong prediction. Two things are important to keep in mind here, however. First, we distinguish between an objective and subjective reading of agentive modals; on an objective reading, if $\lceil S \operatorname{can} \varphi \rceil$ is true, then for some practically available action A, if S tries to do A, she does $\llbracket \varphi \rrbracket^c$; but it may nonetheless be the case that she does not know which action A has this status, and so the ability ascription will not have the *practical effect* of any kind guarantee. Second, as we discuss in a moment, generic ability ascriptions do not guarantee that $\llbracket \varphi \rrbracket^c$ is within S's control, but only something weaker: that $\llbracket \varphi \rrbracket^c$ is in her control if *conditions are normal*. With these caveats in mind, we believe that the characterization of an ability ascription as a kind of hypothetical guarantee does indeed match intuitions, and casts an interesting light on the role of ability ascriptions in practical dialogue.

³⁹This characterization – including its inherent vagueness – helps explain the *impurity* of ability ascriptions, in the sense of Knobe and Szabó (2013); considerations about deontic matters, for instance, influence ability ascriptions insofar as they influence what we count as reasonable and thus practically available.

Our account also answers to an important intuition about what compulsion ascriptions are used to do: when someone says $\lceil I$ cannot but $\varphi \rceil$, she is saying that she will do $\llbracket \varphi \rrbracket^c$ no matter what she tries to do. In other words, compulsion ascriptions serve as a kind of non-hypothetical guarantee. We believe that this matches intuitions about the meaning of claims like this, and their role in practical dialogue.

One upshot of our account is that we predict that ability ascriptions of the form $\lceil S$ can try to $\varphi \rceil$ are quite weak. An ability ascription of this form simply says that there is something practically available such that if S tries to do it, she tries to do $\llbracket \varphi \rrbracket^c$. Thus, for instance, as long as $\llbracket \varphi \rrbracket^c$ is practically available, $\lceil S$ can try to $\varphi \rceil$ will be predicted to be true. This prediction seems right to us. Trying is cheap; anyone can do it.

Does our account make trying too cheap? It predicts that $\lceil S \rangle$ can try to $\varphi \rceil$ is consistent with $\lceil S \rangle$ cannot $\varphi \rceil$, since the latter says that for any practically available action, if $S \rangle$ tries it, she does not do $\llbracket \varphi \rrbracket^c$; and the former says that for some practically available action, if $S \rangle$ tries to do it, she *tries* to do $\llbracket \varphi \rrbracket^c$. This may look wrong; note that (51) is marked:

(51) I can't fix the fridge, but I can try to.

This is not immediately predicted on our account. But we think our prediction is correct: being unable to do something is consistent with being able to try to do it. To see this, note that the third-person variant on (51) in (52) is perfectly coherent:

- (52) John can't fix the fridge; he can only *try* to.
- (53) provides a similar case, using 'able' to select for the agentive reading of the modals, and providing set-up which may make this kind of conjunction more natural:
- (53) Anyone's able to *try* to quit smoking. But few are actually able to do it, and among those, even fewer will. [due to [redacted12] p.c.]

The coherence of (52) and (53) suggests that $\lceil S \text{ cannot } \varphi \rceil$ is indeed consistent with $\lceil S \text{ can try to } \varphi \rceil$. We suspect, then, that the explanation of the infelicity of (51) is pragmatic. The first conjunct of (51) licenses the inference that the speaker believes that, for any practically available action A, if he tries to do A, then he will fail to fix the fridge. And there is something in bad faith about trying to do ψ if you believe that you can't do ψ . Why would you be offering to try, if you knew you would fail? (All this is compatible, of course, with S trying to do ψ in good faith even though S cannot do ψ , as long as S does not *believe* this about herself.) And indeed, note that (51) is improved when this pragmatic constraint is waived, e.g. if the speaker is being pressured to try to fix the fridge. Since no pragmatic story of this kind will extend to third-personal variants of (51), the explanation we've offered makes sense of the contrast between (51) and (52) within our

⁴⁰We are grateful to an anonymous reviewer for this journal for raising this issue, and for suggesting an example like (51), as well as the possibility of a pragmatic response.

framework: both are semantically consistent, as our account predicts, but (51) is (usually) infelicitous for pragmatic reasons.

We note, finally, that tying practical availability to an agent's epistemic state commits us to a certain position on the relation between know-how and ability. We will not offer an analysis of know-how here, but presumably knowing how to do ψ is generally closely connected to knowing of some ρ that doing ρ guarantees doing ψ .⁴¹ On our account, knowing of some ρ that it has this status does not guarantee that one is able to do ψ ; one must also be such that, for some such ρ , if one tries to do ρ , one does ψ . But being able to do ψ does guarantee, *inter alia*, that there is some action ρ that one knows has this status (relative to some description of one's practical situation) — and, moreover, that if one tried to do ρ , one would do ψ . This strikes us as a plausible hypothesis about the relationship between ability ascriptions and know-how: know-how does not guarantee ability, since ability is not a *purely* cognitive matter. But it is a *partly* cognitive matter: ability does entail know-how.⁴²

6 Genericity

The ACA is an account of specific agentive modal ascriptions. In this section we discuss how to generalize the ACA to generic agentive modal ascriptions, like (2): agentive modals whose complement is an action not indexed to a specific time (either explicitly or implicitly).

(2) Mary can touch her nose with her tongue.

We then discuss a variety of related issues involving genericity and agentive modals.

6.1 Generic Ability Ascriptions

We propose that generic ability ascriptions are simply specific ability ascriptions underneath a generic operator. ⁴³ A specific ability ascription has as its prejacent a specific, time-indexed action, like *hitting a bullseye*

⁴¹Not everyone agrees this is always true; see e.g. Brownstein and Michaelson (2015).

⁴²See Stanley (2011) and citations therein for more on this topic.

⁴³See Bhatt (1999), Hacquard (2010), Maier (2015a) for similar ideas, implemented in importantly different ways. For instance, Hacquard gives agentive modals a standard existential quantificational analysis, and then proposes that the fact that their meaning is felt to be stronger than predicted on that analysis stems from an implicit generic operator licensed by imperfective aspect. This approach faces a worry from the fact that the ability modals seem to have a stronger meaning than the existential account would predict *even when they are in contexts which seem not to have a generic operator*. For instance, 'Yesterday at 3 p.m., Susie was able to hit a bullseye' (or its corollary in languages with overt aspect, with 'was able to' in the imperfective) seems to us not to be clearly true, as spoken about Susie the child dart-player, and thus to have a meaning stronger than that predicted by a bare existential analysis. But this sentence does not seem to have an implicit generic operator: generic operators do not happily coexist with temporal phrases like 'yesterday at 3 p.m.' (to see this, note that adding 'generally', which is usually licensed when a covert generic operator is in play, renders this sentence infelicitous; note also that if we replace 'Susie' with an existentially quantified phrase, like 'a child', no generic reading of the existential is available). To handle cases like this, a view along these lines would have to advert to some other, further mechanism of strengthening besides an implicit generic operator. Thanks to Valentine Hacquard for valuable discussion.

in Somerville at 5 pm this Friday, and says of an agent that there is some practically available action such that if she tries to do it, then she does the prejacent. A generic ability ascription has as its prejacent a general action — one which is not indexed to a particular time — like hitting the bullseye, and says of an agent that, in a suitable proportion of normal circumstances, there is an action practically available in those circumstances such that, if she tries to do it, then she does the prejacent.

In the rest of this section we spell out this idea in more detail, though without going too far into the vexed question of what exactly a generic operator is;⁴⁴ we have no particular commitment to the following implementation, which we adopt primarily for concreteness. We assume that genericity is due to an operator GEN — usually tacit, with roughly the meaning of 'generally'. GEN binds a time variable t in a tensed sentence $\psi(t)$ to yield another sentence which says, essentially, that ψ is true in normal *situations* (it doesn't matter for our purposes what a situation is; we could think of it as a world, or part of a world, centered on at least a time, but we assume that sentences can be evaluated for truth not only at a world, but also at situations; and likewise that the conditional selection function f is defined on situations as well as on worlds). More formally, with t a free time variable in $\psi(t)$, t_s the time of a situation s, and 'GEN' subscripted with the variable it binds:

[GEN_t(
$$\psi(t)$$
)] c,w =1 iff in a suitable_c proportion of normal_c situations s , $[\![\psi(t_s)]\!]^{c,s} = 1$.

What counts as a suitable proportion of cases will depend on context (hence the subscript), as will what counts as a normal situation. Roughly, GEN takes an untensed sentence, and returns a sentence that says that the original sentence is *normally* true, according to the standards of normalcy established by the world of evaluation and context.

Two things need to be made explicit in our account of agentive modals before we can show how it interacts with GEN. First, we have left time indices out of our presentation so far for simplicity but we now need to be explicit about them. Ability ascriptions can contain two distinct time indices, one attached to the modal (which says *when* the agent has the ability in question) and one attached to the action (which says when the action takes place). We will ignore the first of these, assuming for simplicity that it does not interact with the GEN operator, and denote the second with a subscript t, so that an ability ascription now has the form of (55), where φ_t denotes a specific action that takes place at t (we change the subject variable to J to avoid confusion with situation variables).

(55) J can φ_t .

Second, we must now represent $A_{S,c}$ as part of the logical form of an ability ascription – i.e. as part of the structure that underlies semantic composition – so that GEN can bind into it. In addition, this parameter

⁴⁴See Krifka et al. (1995, Chapter 1) for an overview, especially §1.2.6 on the approach we follow. See Sterken (2015), Nickel (2016) and citations therein for more recent discussion.

will now have a time argument, which represents the time at which the actions in question are available, and which - again, perhaps simplifying - we assume is the same as the index on the action. GEN binds both these time indices. On our account generic ability ascriptions thus have the form and truth conditions in (56):

(56)
$$[\![GEN_t(J \text{ can } (\mathcal{A}_{J,c,t})(\varphi_t))]\!]^{c,w} = 1 \text{ iff at a suitable}_c \text{ proportion of normal}_c \text{ situations } s: \exists A \in \mathcal{A}_{J,c,t_s}: [\![\varphi_{t_s}(J)]\!]^{c,f_c(s,try(J,A))} = 1.$$

Informally, a generic ability ascription with the surface form $\lceil J \operatorname{can} \varphi \rceil$ says of J that, in enough normal situations s, there is a practically available action at the time of s which is such that if J tries to take that action, J brings it about that $\llbracket \varphi \rrbracket^c$ holds of her at that time. In short: in normal situations, J can do $\llbracket \varphi \rrbracket^c$.

To make these truth-conditions plausible, we must assume that what counts as a 'normal circumstance' depends on context as well as the prejacent. For example, consider (57):

(57) George can make great ratatouille.

When evaluating (57), the set of normal situations will include only those where George has all the ingredients, tools, etc. needed to make ratatouille. Note that this kind of sensitivity is needed for generics in general; for instance, we must say something similar in evaluating (58):

(58) George makes great ratatouille.

Our account gives intuitive truth conditions for generic ability ascriptions. It also makes an interesting further prediction: the possibility of scopal interactions between GEN and other constituents of ability ascriptions. We believe this prediction is borne out, and explains a well-known ambiguity in ability ascriptions. ⁴⁵ Consider the following two possible logical forms for a generic ability ascription:

- (59) GEN_t(J can $(A_{J,c,t})(\varphi_t)$)
- (60) J can $(A_{J,c,t})$ (GEN_t (φ_t))

We've assumed that generic ability ascriptions have the form of (59), and mean that, in enough normal situations, there is something available for J to do such that if she tries to do it, she does $[\![\varphi]\!]^c$. But if GEN can scope under $\mathcal{A}_{J,c,t}$, as in (60), we will find generic ability ascriptions which mean that there is something (now) available for J to do such that if she tries to do it, then in enough normal circumstances, she will do $[\![\varphi]\!]^c$. A6 Both of these two scopal possibilities seem to be manifested. Consider (61):

(61) John can speak Greek.

⁴⁵We are indebted to Justin Khoo for first suggesting this possibility to us.

⁴⁶Assuming that free time variables are set by default to the time of utterance.

As Lewis (1976) observes (broadly following Aristotle (1968, 417a)), a sentence like (61) can mean two things. First, it can mean that John is able, in general, to speak Greek, i.e. is a fluent speaker of Greek. Second, it can mean that John has the physiological and cognitive apparatus to come to be a fluent speaker of Greek, by contrast, say, with an ape; John could learn Greek if he wanted to. These two interpretations correspond precisely to the two logical forms (59) and (60), respectively. The first parse says that, in most normal situations, there is something practically available which John can try to do in order to speak Greek. The second parse says there is something practically available such that if John tries to do it, then, in normal circumstances, he speaks Greek. In short, the first says there is normally something John can do to speak Greek; the second says there is something John can do to become a normal Greek-speaker. The possibility of scopal interactions which the present account predicts and explains this ambiguity in ability ascriptions.⁴⁷

6.2 Uncertain Abilities

As we noted in Section 5.2, we often ascribe abilities to agents even when they are not certain to succeed at a given action should they try. Cases like this bring up a number of interesting issues. We touched on this topic in discussing the golfer Jones, where we said that (29) can have a true reading even if Jones actually misses, provided there was some knowledge that Jones lacked such that, if he had had it, he would have made the shot.

(29) Jones is able to make this shot right now.

But not all uncertain ability ascriptions can be handled by an appeal to lack of knowledge along these lines. In this section, we show that by carefully distinguishing generic from specific agentive modal claims — as well as objective from subjective ones — the ACA gives a satisfying account of these *uncertain ability ascriptions*.

Suppose Stephen Curry has a difficult three-point shot to make. Consider (62):

(62) Curry can make this shot.

Intuitively, given Curry's skill, (62) is true in many contexts, even though in many cases — perhaps the majority — Curry will miss when he tries to make a shot like this. How do we predict this? In light of the present discussion, we distinguish two readings of (62): a generic reading and a specific reading. On the generic reading, (62) means something like (63):

(63) He can make this *kind* of shot.

⁴⁷This explanation of the ambiguity presupposes that it arises only for generic, not specific, ability ascriptions. This, however, seems to match intuitions.

Even if Curry was not able to make the very shot he tried to make, we can predict that sentences like (63) are true, if Curry is able to make this kind of shot in a suitable proportion of *normal cases*. What counts as a suitable proportion of normal cases is, again, highly context-dependent. In a case where the baseline of successful shots is quite low, what counts as a suitable proportion may likewise be quite low. Compare a generic like 'Rats carry bubonic plague'. This may be true even though only a small proportion of rats in fact carry the plague.⁴⁸ Thus in a situation that involves making a highly risky shot, it may suffice for the truth of a generic like (63) that Curry succeeds half the time. A generic interpretation of (62) thus may account for many true readings of uncertain ability ascriptions like (62).

What about specific readings of (62), forced by adding an overt time index, as in (64)?

(64) Curry can make this shot right now.

Our predictions about this case will depend on how we flesh it out. We distinguish cases in which Curry is ignorant in a relevant way about his practical situation (*ignorance cases*) from those in which he has no relevant ignorance (*no-ignorance cases*).

Begin with ignorance cases. We can treat these cases just as we treated Jones' golf shot above: in these cases, (64) will have a true objective reading (but not a true subjective reading). Imagine, for example, that Curry did not know that his teammate was going to set a screen for him. Had he known about the screen, he would have used it, and made the shot; in fact, he didn't know about it, and so he went for a different shot and missed. Then we predict a true, objective reading of (64), which might be elicited by reflection like (65):

- (65) I was able to make that shot; the only reason I didn't is that I didn't know about the screen.
- (64) will also have a false, subjective reading; after all, Curry didn't know about the screen.

Consider now no-ignorance cases — cases in which Curry is not ignorant about any relevant features of his practical situation. These subdivide into two classes. One kind of case involves a failure of execution. Imagine, for example, that Curry tried to do a specific action which he knew was a way to make the shot, but his body failed him — he was tired and so was not quick enough to escape his defender. In a case like this, our intuition is that Curry was simply not able to make that very shot — he was just too tired to execute the necessary moves. Assuming there is no other action such that, had he tried to do it, he would have made the shot, then we rightly predict that (64) has no true reading in this scenario.

The second kind of no-ignorance case is one in which Curry's body does not fail him, but for some reason he nonetheless fails to execute the specific game-winning action. We can imagine this happening in cases of inattention or lack of focus. In this case, we are inclined to think there is a true reading of (64):

⁴⁸See again Krifka et al. (1995) and citations therein.

Curry was able to make the shot; he just had to pay more attention. Suppose that Curry knew that A was a way to make the shot, given what he knew about his practical situation. And suppose that Curry tried to make the shot, but failed. We can still predict a true reading of (64), if we maintain that the action Curry tried was not A. A may be a highly specific and taxing action to attempt (aiming with this arc and spin from this position); Curry may have simply failed to try to do that action (instead simply trying something non-specific like making a shot). A lack of focus, in other words, may involve a failure to try to do the requisite very specific action. (A common complaint from coaches: 'You're not really trying!'). If that's the right diagnosis, we predict a true subjective reading of (64): A is a practically available action for Curry, and if Curry had tried to do A, he would have made the shot. He just didn't try to do A.

Uncertain ability ascriptions like those discussed here bring up a variety of complex issues; our distinctions between objective versus subjective — and generic versus specific — agentive modals help make sense of the subtle intuitions in these cases.

6.3 Kenny's Objection

Our account of genericity helps answer a classic objection to modal analyses of ability, due to Kenny (1976). Kenny points out that on a modal analysis like the orthodox one, $\lceil S \operatorname{can} (\varphi \operatorname{or} \psi) \rceil$ is predicted to entail $\lceil (S \operatorname{can} \varphi) \operatorname{or} (S \operatorname{can} \psi) \rceil$. This entailment also follows from the ACA. For suppose that (66) is true:

(66) S can
$$(\varphi \text{ or } \psi)$$
.

Then there is a practically available action A such that the closest world where S tries to do A is one where she does $[\![\varphi]\!]^c$. That world is either one where she does $[\![\varphi]\!]^c$ or one where she does $[\![\psi]\!]^c$. It follows from our semantics that, if we hold the context fixed, (67) is true.

(67) (S can
$$\varphi$$
) or (S can ψ).

But this entailment conflicts with intuition. (68) does not obviously entail (69):

- (68) Leo can, right now, draw a red card or draw a black.
- (69) Leo can, right now, draw a red card or Leo can, right now, draw a black card.

We have two things to say in response. The first response says nothing about genericity. It parallels the response of a defender of conditional excluded middle to purported counterexamples. Focus on the specific readings of (67) and (68). We hold that a sentence like (68) is true because it is determinately true that there is something practically available such that if Leo tries to do it, he draws a red card or draws a black card (namely, drawing a card at random); but it is not determinately true that Leo can right now draw a red card; nor is it determinately true that Leo can right now draw a black card. Nonetheless, the disjunction (69) is

determinately true, since no matter which way we fill things out, one of these disjuncts is true — even though we don't know, and there may be no fact of the matter, which. We think this response makes good sense of our intuitions about the relation between (68) and (69). When we reflect, we *are* inclined to think that (69) follows from (68), even though we don't know — and there may be no fact of the matter — which of 'Leo can draw a red card' or 'Leo can draw a black card' is true.

The second response to Kenny's objection appeals to genericity. We predict that the entailment from (66) to (67) will be *invalid* when the ability ascriptions in question are read generically, since (70) obviously does not entail (71):

- (70) GEN_t(J can $(A_{J,c,t})(\varphi_t \vee \psi_t)$)
- (71) $\operatorname{GEN}_t(\operatorname{J}\operatorname{can}(\mathcal{A}_{J,c,t})(\varphi_t))\vee\operatorname{GEN}_t(\operatorname{J}\operatorname{can}(\mathcal{A}_{J,c,t})(\psi_t))$

It might be the case that, in a suitable proportion of normal situations, J can do $[\![\varphi \lor \psi]\!]^c$, without it being the case that, in a suitable proportion of normal situations, J can do $[\![\varphi]\!]^c$, or that, in a suitable proportion of normal situations, J can do $[\![\psi]\!]^c$. We thus straightforwardly predict that (72) fails to entail (73):

- (72) Leo can (generally) draw a red card, or draw a black card.
- (73) Leo can (generally) draw a red card, or Leo can (generally) draw a black card.

6.4 Masks and Finks

Our distinction between specific and generic abilities also helps us respond to another classic objection to the CA, an objection which also extends to our account: namely, cases involving *masks* and *finks* (see e.g. Martin (1994), Fara (2008)). Clearly distinguishing between specific and generic ability ascriptions allows us to dismiss these objections.

Suppose I am able to break the vase on my desk if I drop it. But suppose the vase is wrapped in bubble wrap, which I can't presently remove. The bubble wrap *masks* my ability to break the vase. So I am able to break the vase, even though there is no practically available action such that if I try to do it, I break the vase. What should we say about this case?

In this case, we maintain that the *specific* ability ascription (74) is, in fact, false.

(74) I can break this vase right now by dropping it.

I *can't* break the vase right now; the bubble wrap would thwart my attempt to do so. But the mask leaves intact my *generic* ability to break the vase. Consider (75):

(75) I can break the vase.

(75) is true in the case described — under normal circumstances the vase would break if I dropped it, since it would not be wrapped in bubble wrap. I have a *generic* ability to break this vase but do not have a specific ability to break the vase *right now*, since it is wrapped in bubble wrap. And if I do not have a specific ability to break the vase, this case poses no counterexample to our account.

We can respond similarly to concerns about finks. Once we clearly distinguish specific from generic ability ascriptions, we see that finks and masks do not threaten our account.

6.5 Trying

Not everything that can be the grammatical subject of an agentive modal ascription can try. This is a *prima facie* problem for our account. We close this section by showing how genericity can help with this problem.

On the one hand, we think the notion of trying extends further than one might first imagine; consider for instance (76):

(76) The flower can follow the sun.

[due to Kieran Setiya p.c.]

We think it plausible that we are conceiving of the flower as *trying* to follow the sun; after all, we might say that the flower is trying to get as much light as possible. But in other cases it is not plausible to think of the subject of an ability ascription as trying, as in (77):

(77) This elevator is able to carry 3,000 pounds. [due to Due to Irene Heim (p.c.),/ attributed to Maria Bittner, p.c.]

It doesn't make sense to say that the elevator *tries* to carry 3,000 pounds. We can dramatize this point with examples like (78):

(78) This rock can tip the balance.

[due to Martin Hackl, p.c.]

Rocks do not try to do things. Furthermore, these modals look like ability modals, not circumstantial modals; they can be paraphrased as 'has the ability to'.

How can we accommodate these examples? Here's a sketch of a response. (77) can be glossed as saying, roughly, that there is some available action such that in the closest world where *the operator of the elevator* tries to take that action (in this case, loading the elevator with 3,000 pounds of freight and pressing 'up'), the elevator carries 3,000 pounds. Likewise, (78) can be glossed as saying, roughly, that there is some available action such that in the closest world where *some relevant agent* tries to take that action (in this case, putting the rock on the balance), the rock tips the balance.

More formally, we say that GEN can bind into the *subject index* on the set of actions and in the selection function, in addition to the time index. Ignoring tense, a sentence like (77) would thus have the form and

meaning:

(79)
$$[\![\operatorname{GEN}_R(\operatorname{J}\operatorname{can}(\mathcal{A}_{R,c})(\varphi))]\!]^{c,w} = 1 \text{ iff}$$
 for a suitable proportion of normal agents $R: \exists A \in \mathcal{A}_{R,c}: [\![\varphi(J)]\!]^{c,f_c(w,try(R,A))} = 1$

Informally: enough normal agents are such that, in normal situations, there is some practically available action A such that if they try to do A, they bring it about that J does φ . Thus, e.g., (77) will mean that most normal agents can make it the case that the elevator carries 3,000 pounds right now; (78) will mean that most normal agents can make it the case that the rock tips the balance. The ACA can thus accommodate agentive modals whose grammatical subject is not an agent in any intuitive sense by appealing to genericity.

Note that a similar move may be possible in order to account for some (though certainly not all) '-ble' and '-ile' constructions, like 'breakable'. Building constructions like these out of ability ascriptions may look implausible; after all, (80) clearly does not say anything about whether the vase would break *if it tried something*.

(80) This vase is breakable.

But we may posit an implicit generic agent who does the relevant trying, so that (80) means that, in normal circumstances, a normal agent can break the vase. On this approach, (80) has a generic operator that binds an implicit subject argument. -ble and -ile constructions are complex and varied, but this may provide a basis for an analysis of the subset of these constructions which are, in some broad sense, about ability.⁴⁹

7 Conclusion

We have advocated a new semantics for agentive modals, the *act conditional analysis*. Our semantics builds on an old idea about ability: that what one can do depends on what would happen if one tried. But our account avoids the problems that led us, and most other researchers, to reject the conditional analysis. We achieve this by incorporating quantification over a contextually supplied set of actions into the meaning of 'can'. This approach makes plausible predictions about agentive modals, avoiding the array of objections we enumerated for the conditional analysis, as well as the orthodox and universal approaches.

Our account highlights the relation between ability modals and compulsion modals. We have argued that making the right predictions about the meaning of compulsion modals provides a crucial criterion of adequacy for accounts of ability modals, and have suggested that distinguishing the class of compulsion modals — and the broader class of agentive modals, which includes both ability and compulsion modals — is critical for the theory of modals in natural language, as well as for an array of broader philosophical

⁴⁹See Maier (2015c) and Vetter (2014) for recent discussions related to the present proposal.

issues.

We conclude by mentioning some questions that we have left open. First, we think that compulsion modals may play a role in making sense of intuitions about freedom and compulsion in the philosophical debate on free will. Appeal to the notion of practical availability may allow us to make sense of the malleable and contradictory intuitions in that debate, as well as to make precise the traditional idea that freedom, on the one hand, and ability and compulsion, on the other, are closely connected.

Another important question which we leave open is the cross-modal picture which our account suggests. On our account, agentive 'can' differs in important ways from circumstantial, deontic, and epistemic 'can', at least as those are treated in the orthodox framework. Like the orthodoxy, our semantics treats agentive 'can' as an existential quantifier; but in contrast to the orthodoxy, in our account, 'can' quantifies over a set of actions (i.e. properties), not a set of worlds; and a conditional selection function plays a crucial role in its semantics. Thus, if we leave unchanged the orthodox account of circumstantial, deontic, and epistemic 'can', we must maintain that 'can' is lexically polysemous: it can be interpreted either as an ordinary existential quantifier, to yield a circumstantial, deontic, or epistemic reading; or it can be interpreted as on the ACA, to yield an agentive reading.

Is this proposed polysemy a substantial drawback? The answer to this will depend partly on the cross-linguistic picture — how uniformly languages use the same word for an agentive modal which they use for other kinds of modals. We should note, though, that we can in theory preserve the predictions of the ACA, while preserving the uniformity of the orthodox framework, by shoehorning the ACA into the orthodox framework. For any agent S, just let the modal base take any world w to \emptyset and let the ordering source take any world w to $\{w'\}: \exists A \in \mathcal{A}_{S,c}(w'=f_c(w,try(S,A))\}$. This account makes the same predictions as ours.

But — assuming that the elements of the semantic formalism are meant to map onto a meta-semantic story in some straightforward way — it involves awkward meta-semantic commitments. A meta-semantics for both the ACA and this orthodox recasting of the ACA must account for how $\mathcal{A}_{S,c}$ and f_c are determined, since both views refer to these parameters. But the orthodox recasting *also* makes reference to a highly gerrymandered ordering source. It is hard to see how this ordering source could play any role in an explanatory meta-semantics. So we think there is some reason simply to dispense with it, and to adopt the ACA in the form in which we have spelled it out.

Unification, however, might also be possible in a different direction: not by shoehorning the ACA into the orthodox framework, but rather by replacing the orthodox framework with a unified framework structured more like the ACA. Indeed, a number of recent proposals have departed from orthodoxy in a manner similar to the ACA, incorporating quantification over sets of properties or propositions into the meaning of modals.

This raises the possibility that our proposal could be part of a new, unified approach to modality.⁵⁰ We leave further exploration of this question for future work - a question which will turn, again, on further work on the semantics, meta-semantics, and cross-linguistic typology of modals.

A final important question which we leave open is how ability modals interact with tense and aspect to yield 'actuality implications' from $\lceil S \rceil$ was able to $\varphi \rceil$ to $\lceil S \rceil$ did $\varphi \rceil$ (see e.g. Bhatt (1999), Hacquard (2006)).

There is more to say about agentive modals, then. However these details are worked out, we hope to advance the literature by laying the framework for an account which answers to an important intuition about what agentive modals are used to do. When someone says $\lceil I \operatorname{can} \varphi \rceil$, she is giving a kind of hypothetical guarantee about her relationship to $\llbracket \varphi \rrbracket^c$. She is assuring her interlocutors that $\llbracket \varphi \rrbracket^c$ is within her control in a certain way — that is, that there is some action that she could reasonably conclude in favor of doing, such that if she tries to do it, she succeeds in doing $\llbracket \varphi \rrbracket^c$. And when someone says $\lceil I \operatorname{cannot} \operatorname{but} \varphi \rceil$, she is giving a kind of non-hypothetical guarantee about her relationship to $\llbracket \varphi \rrbracket^c$. She is assuring her interlocutors that refraining from $\llbracket \varphi \rrbracket^c$ is not an option for her; that is, that she will $\llbracket \varphi \rrbracket^c$ no matter what.

References

Albritton, R. (1985). Freedom of will and freedom of action. In *Proceedings and Addresses of the American Philosophical Association*, volume 59, pages 239–251. American Philosophical Association.

Aristotle (1968). De anima. Clarendon Press.

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

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

Belnap, N. (1991). Backwards and forwards in the modal logic of agency. Philosophy and Phenomenological Research, 51(4):777-807.

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.

Brennan, V. M. (1993). Root and epistemic modal auxiliary verbs. PhD thesis, University of Massachusetts Amherst.

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

Brownstein, M. and Michaelson, E. (2015). Doing without believing: Intellectualism, knowledge-how, and belief-attribution. *Synthese*, pages 1–22. 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. Chierchia, G. (1989). Anaphora and attitudes *De Se*. In Bartsch, R., van Benthem, J., and van Emde Boas, P., editors, *Semantics and Contextual Expression*, pages 1–31. Dordrecht Reidel.

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.

Fara, M. (2008). Masked abilities and compatibilism. Mind, 117(468):843-865.

von Fintel, K. (1997). Bare plurals, bare conditionals, and Only. Journal of Semantics, 14:1-56.

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.

Hacquard, V. (2006). Aspects of Modality. PhD thesis, MIT.

Hacquard, V. (2010). Ability modals: A ♦ in the generic rough. Slides from the Dispositions, Abilities and States Workshop.

Hawthorne, J. (2005). Chance and counteractuals. Philosophy and Phenomenological Research, 70(2):396-405.

Hilpinen, R. (1969). An analysis of relativised modalities. In Davis, J., Hockney, D., and Wilson, W., editors, *Philosophical Logic*, volume 20. Springer.

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

Kamp, H. (1973). Free choice permission. In Proceedings of the Aristotelian Society, volume 74, pages 57-74.

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

⁵⁰For recent work along structurally similar lines, see e.g. Cariani et al. (2013) and Cariani (2013) for related approaches to deontic modals; Moss (2015) on epistemic modals; Yalcin (2012) and Willer (2013) on attitude verbs; and Villalta (2008) on verbs of desire.

Knobe, J. and Szabó, Z. G. (2013). Modals with a taste of the deontic. Semantics and Pragmatics, 6(1):1–42.

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.

Krifka, M., Pelletier, F. J., Carlson, G. N., ter Meulen, A., Link, G., and Chierchia, G. (1995). Genericity: An introduction. In Carlson, G. N. and Pelletier, F. J., editors, *The Generics Book*, pages 1–124. University of Chicago Press, Chicago.

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.

Lewis, D. (1973). Counterfactuals. Harvard University Press, Cambridge, MA.

Lewis, D. (1976). The paradoxes of time travel. American Philosophical Quarterly, (145-152).

Lewis, D. (1981). Ordering semantics and premise semantics for counterfactuals. Journal of Philosophical Logic, 10(2):217-234.

Lewis, D. (1996). Elusive knowledge. Australasian Journal of Philosophy, 74(4):549-567.

Maier, J. (2015a). Ability, modality, and genericity. Manuscript.

Maier, J. (2015b). The agentive modalities. Philosophy and Phenomenological Research, 40(1):113-134.

Maier, J. (2015c). Dispositions and ergativity. The Philosophical Quarterly, 65(260):381-395.

Martin, C. (1994). Dispositions and conditionals. The Philosophical Quarterly, 44:1-8.

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

Moss, S. (2015). On the semantics and pragmatics of epistemic vocabulary. Semantics and Pragmatics, 8:1-81.

Nickel, B. (2016). Between Logic and the World: An Integrated Theory of Generics. Oxford University Press.

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

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

Sharvit, Y. (2003). Trying to be progressive: the extensionality of try. *Journal of Semantics*, 20:403–445.

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

Stalnaker, R. (1975). Indicative conditionals. Philosophia, 5(3):269-86.

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.

Stanley, J. (2011). Know-How. Oxford University Press.

Sterken, R. K. (2015). Generics in context. Philosophers' Imprint, 15(21):1-30.

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

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

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

Vetter, B. (2014). Dispositions without conditionals. Mind, 123(489):129-156.

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.

Woodward, A. L. (1998). Infants selectively encode the goal object of an actor's reach. Cognition, 69(1):1-34.

Wurmbrand, S. (2002). Syntactic vs. semantic control. In Zwart, J.-W. and Abraham, W., editors, *Studies in Comparative Germanic Syntax*, pages 95–129.

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