

Two Puzzles about Ability *Can**

Malte Willer

Abstract

The received wisdom on ability modals is that they differ from their epistemic and deontic cousins in what inferences they license and better receive a universal or conditional analysis instead of an existential one. The goal of this paper is to sharpen the empirical picture about the semantics of ability modals, and to propose an analysis that explains what makes the *can* of ability so special but that also preserves the crucial idea that all uses of *can* share a common lexical semantics. The resulting framework combines tools and techniques from dynamic and inquisitive semantics with insights from the literature of the the role of agency in deontic logic. It explains not only why the *can* of ability, while essentially being an existential modal operator, sometimes resists distribution over disjunction and interacts with its duals in particular and hitherto unnoticed ways, but also has a tendency to license free choice inferences.

1 The Plot

Ability modals—modals that are used to state what someone *can* do—differ from other modals in interesting ways. The initial observation goes back to Kenny (1975, 1976), who notes that the *can* of ability seems to resist the inference rule of distribution over disjunction. Imagine that Mary is skilled enough to hit the board, but not skilled enough to hit any particular region of the board. Against this background, it seems as if there is a reading of *can* on which we would accept (1a) but not accept (1b) or (1c):

- (1) a. Mary can hit the board.
- b. Mary can hit the top half of the board.
- c. Mary can hit the lower half of the board.

But to hit the board is to hit its top or its lower half and thus, so the story continues, we have a case in which what looks like an existential modal operator does not distribute over disjunction: a sentence that is apparently equivalent to a disjunctive possibility $\lceil \diamond(\phi \vee \psi) \rceil$ fails to entail $\lceil \diamond\phi \vee \diamond\psi \rceil$ (see also Horty 2001 and Portner 2009). This marks

*Forthcoming in *Linguistics and Philosophy*. Thanks to three anonymous L&P referees as well as to audiences at the workshop on Inquisitiveness Below and Beyond the Sentence Boundary, the University of Chicago, the University of Dubrovnik, the 2018 Central APA, and the Martin-Luther-University of Halle-Wittenberg for comments and discussion. David Boylan and Paul Portner provided detailed and constructive written comments that proved to be invaluable in improving this manuscript. Special thanks also to Melissa Fusco, Jeroen Groenendijk, Jeff Horty, Johannes Hübner, Matt Mandelkern, and Floris Roelofsen for illuminating discussions. The usual disclaimers apply.

a striking contrast between ability *can* and run-of-the-mill existential modals such as epistemic *might* and deontic *may*, which happily distribute over disjunction. To wit, it seems plain weird to accept (2a) without accepting (2b) or (2c), and no less weird to accept (3a) without accepting (3b) or (3c):

- (2) a. Mary might hit the board.
- b. Mary might hit the top half of the board.
- c. Mary might hit the lower half of the board.
- (3) a. Mary may/is permitted to hit the board.
- b. Mary may/is permitted to hit the top half of the board.
- c. Mary may/is permitted to hit the lower half of the board.

We thus need an explanation for why the *can* of ability in particular seems to resist distribution over disjunction, while ordinary possibility modals—such as epistemic and deontic ones—do not.

Kenny’s own conclusion is that ability *can* better not receive a possible worlds analysis in the first place. Few have drawn this specific moral, but the basic message that natural language “can” resists a uniform semantic analysis remains dominant until today.¹ One prominent view is that the *can* of ability patterns with epistemic and deontic *must* in that it is a universal quantifier over a set of points of evaluation (see Brown 1988; Giannakidou 2001; and Giannakidou and Staraki 2012).² Another is that the *can* of ability is best understood as a conditional operator stating, roughly, what would happen if the subject tried to bring about the prejacent (see Cross 1986; Mandelkern et al. 2017; and Thomason 2005). What all of these accounts have in common—without further maneuvers anyway—is that the semantic contribution of *can* in (4a) is different in kind from the one in (4b) and (4c):

- (4) a. Mary can swim.
- b. You can go now.
- c. There can be life on mars.

On the views under consideration, *can* denotes a universal or conditional operator in an ability attribution, as in (4a). On its deontic interpretation in (4b), *can* is existential, and likewise on its epistemic interpretation in (4c).

Everybody, I think, can agree that giving up on a uniform analysis of *can* comes with theoretical costs. As Kratzer (1977, 1981, 1991, 2012) remarks, an ideal theory would assign to all uses of *can* and *must* in discourse a common semantic core while allowing for differences in modal flavor (epistemic, deontic, agentive, and so on). A semantics that essentially provides a separate semantic entry for ability *can* falls short of this desideratum, and with substantial repercussions down the line. Note, for instance, that there are several instances in which the *can* of ability does not only distribute over disjunction but in fact seems to behave like other possibility modals in that it licenses the even stronger *free choice* inference. Consider:

¹Hackl’s (1998) analysis is a notable exception, but it does not address Kenny’s objection.

²Kenny (1975, p. 139) briefly floats this idea without endorsing it. Portner (2009) proposes analyzing ability *can* as expressing a “good” possibility in the sense of Kratzer 1981, 1991, thus effectively combining existential with universal quantification. This is also the key idea behind Horty’s (2001) analysis, which is articulated in stit-semantics (cf. Belnap et al. 2001).

- (5) We can speak Dutch or French with each other.

An utterance of (5) strongly suggests that we can speak Dutch with each other, *and* that we can speak French with each other.³ This matters for current purposes since disjunction scoping under existential modals such as *may* and *might* reliably gives rise to the free choice effect:⁴

- (6) They might be speaking Dutch or French with each other.
a. \rightsquigarrow They might be speaking Dutch with each other.
b. \rightsquigarrow They might be speaking French with each other.
- (7) We may speak Dutch or French with each other.
a. \rightsquigarrow We may speak Dutch with each other.
b. \rightsquigarrow We may speak French with each other.

There is every reason to insist that whatever mechanism explains why the free choice effect arises for *might* and *may* should also apply, at least in principle, to *can*—otherwise we would have to give two separate accounts of what looks like one and the same basic phenomenon. The most promising strategy is to assign to the modals at play a common lexical semantics that predicts free choice, perhaps in combination with plausible pragmatic principles.⁵ But this path is not, or at least not immediately, available if *might* and *may* are existential in nature while agentive *can* is universal or conditional.

Nothing said so far is meant to suggest that there is nothing special about ability *can*. Its resistance to distribution over disjunction, in fact, does not mark the only contrast with epistemic *might* and deontic *may*. Here is another. Mandelkern et al. (2017) observe that ability *can* seems to allow for duals, which they label “compulsion modals” and occur in the following examples:

- (8) a. Lara cannot but eat another cookie right now.
b. I have to sneeze right now.
c. I cannot *not* eat another cookie.

³Nouwen (forthcoming) suggests that free choice for ability modals is a limited phenomenon: it is crucial to the example that it involves two variations of the same ability. Specifically, Nouwen suggests that the following sentence does not have a free choice reading since the preadjacent refers to two homogeneous abilities:

Betty can balance a fishing rod on her nose or juggle four hot potatoes with just her left hand.

It strikes me as uncontroversial, however, that a free choice reading can easily be made salient. For instance, if the sentence is uttered in response to the question of why we should hire Betty as a performer at our upcoming party, it clearly suggests that her portfolio includes both of these amazing tricks. One might, of course, insist that in this context the balancing and the juggling are variations of what Betty can do at our party. The obvious worry then is that the notion of what counts as a variation of the same ability is too fluid to impose a significant restriction on the empirical distribution of free choice abilities.

⁴Kamp’s (1973, 1978) discussion of the free choice effect is seminal, though the label goes back to von Wright (1968).

⁵This strategy dominates the literature. Pragmatic implementations include Alonso-Ovalle 2006; Fox 2007; Franke 2011; Klinedinst 2007; Kratzer and Shimoyama 2002; and Schulz 2005. Semantic implementations include Aher 2012; Aloni 2007, ms.; Barker 2010; Fusco (2015a, 2015b); Geurts 2005; Goldstein forthcoming; Simons 2005; Starr 2016; Willer 2015, 2017, 2018; and Zimmermann 2000.

It is reasonable to demand that the meaning of compulsion modals is derivable in the same way in which we derive the meaning of other necessity modals, namely as the dual of the relevant possibility modal. As Mandelkern et al. (2017) detail, this requirement is difficult to satisfy for a range of traditional analyses of ability *can*,⁶ but here I would like to draw attention to another surprising fact, namely that the negation of ability *can* does not trivially entail the corresponding compulsion statement. Consider:

- (9) a. Bob can't hit the bullseye.
- b. Bob cannot but not hit the bullseye
- c. Bob can't play Liszt's *Campanella*.
- d. Bob cannot but not play Liszt's *Campanella*.

Insofar as (9a) and (9c) say that Bob fails to be in a position to perform the actions described by the prejacent, their truth does not entail (9b) or (9d), respectively. The latter seem to say that Bob's actions inevitably aim at avoiding to hit the bullseye or to play Liszt's *Campanella*, and this is different from just not being in a position to hit a certain small spot on the board or to perform a notoriously difficult composition.

The previous observation is in principle not too surprising, since the absence of an ability to do something does not guarantee the presence of a compulsion not to do it—no wonder that negated *cans* do not trivially entail their corresponding compulsion statements. It does mark, however, another contrast with epistemic *might* and deontic *may*. Consider:

- (10) a. Bob may not/is not permitted to hit the bullseye.
- b. Bob must not/is required not to hit the bullseye.
- c. Bob can't be playing Liszt's *Campanella*.
- d. Bob must not be playing Liszt's *Campanella*.

(10a) and (10b) as well as (10c) and (10d) seem to be equivalent. So when it comes to epistemic and deontic modals, the negation of a possibility does seem to entail the necessity of the corresponding negation. The way deontic and epistemic modals play with negation thus differs from what we have observed about ability *can*.

There is thus every reason to think that the *can* of ability differs from run-of-the-mill existential modals in interesting ways. The point remains that we also have every reason to aim for a semantic analysis that preserves a distinct sense of uniformity between the *can* of ability and its deontic and epistemic incarnations. The goal of this paper is to demonstrate that such a story can be told. Its key claim, in brief, is that *can* requires possibility across the board: what makes the *can* of ability special is that it requires the possibility of a certain action, while epistemic and deontic *can* require the possibility of a certain state of affairs. This proposal allows us to explain the differences between ability

⁶To illustrate the nature of the problem, consider Brown's (1988) analysis, on which "Mary can hit the board" (which is of the form ' $\diamond\phi$ ') is true just in case there is an action available that guarantees the outcome of her hitting the board. Then "Mary cannot but hit the board" (which, assuming that "can" and "cannot but" are duals, is of the form ' $\neg\diamond\neg\phi$ ') is true just in case no action available guarantees that she fails to hit the board, that is, just in case every action available leaves the possibility of her hitting the board open. But that just seems to get the truth-conditions of the latter sentence all wrong: "Mary cannot but hit the board" seems to say that Mary will hit the board no matter what she does, not that she *might* hit the board no matter what she does.

can and its deontic and epistemic cousins in a familiar fashion: in terms of differences between the relevant modal domains.

I begin with an informal outline of the key ideas and concepts of the proposal (Section 2) and then dive into the details (Section 3). Section 4 addresses a few remaining issues. Section 5 concludes the discussion.

2 Basics

It is a familiar idea that understanding ability *can* requires some conception of agency: to say that Mary can hit the board is, after all, to say that Mary is in a position *to do* something (see Hackl 1998 and Horty 2001, among many others). In fact, there is good reason to think that the role of agency is key to generating our puzzles about distribution over disjunction and negation. Go back to Kenny’s scenario, for instance, and suppose that we shift our attention away from Mary and toward her dart, which is a mere participant in the dart throwing event. In this case, distribution over disjunction seems entirely unproblematic, and so does the interaction between *can* and its dual *cannot* *but*:

- (11) a. The dart can hit the board.
b. The dart can hit the top half of the board.
c. The dart can hit the lower half of the board.
- (12) a. The dart can’t hit the board.
b. The dart cannot but not hit the board.

There is no inclination to accept (11a) without accepting (11b) or (11c), and (12a) entails (12b) without further ado. It thus matters for the logic of *can* whether the subject is a genuine actor or merely participates in an event: only if the former is the case does *can* behave in the very special ways we observed.

The obvious question now is what gives ability *can* its distinct agentive flavor and how (if at all) this helps us understand its special meaning. The answer I propose here is that ability *can* has an agentive flavor because it requires that it be possible for its subject to perform a certain action—the action described by the prejacent.⁷ And I propose that we can make reasonable assumptions about how to model actions that, if taken together, make sense of the data observed in the previous section. The first one is this: actions are (for current purposes anyway) well-understood in terms of their possible outcomes—the ways the world could be that are compatible with implementing that action. Mary can hit the board, for instance, just in case it is possible for her to perform the action of hitting the board, and that is so just in case her action space includes a proposition—a set of possible worlds—which entails that Mary hits the board. This assumption will eventually allow us to account for distribution over disjunction failures for ability *can*.

The second assumption is this: we will also allow for the possible worlds used in modeling actions to be *partial* in order to account for the fact that not all omissions are created equal, and in particular for the distinction between *failing* to do something and

⁷The spirit of this proposal is sympathetic to Maier’s (2013) claim that ability *can* attributes options, though the specific line of inquiry he pursues is different from mine.

refraining from doing something.⁸ Failing to hit the board and refraining from hitting the board are, for instance, not the same actions. The former is compatible with trying to hit the board, while the latter actually requires the presence of some intention that leads to a behavior that is incompatible with hitting the board (see e.g. Clarke 2014 for detailed discussion). Accordingly, we distinguish in principle between an individual’s failing to be in the extension of a predicate and his or her falling into the anti-extension of that predicate, and we will take the latter to mean that the individual refrains from performing—and not just fails to perform—the activity denoted by the predicate. This assumption will eventually allow us to account for the fact that the negation of ability *can* does not trivially entail the corresponding compulsion statement.

I have appealed to the notion of an available action, and I add that this notion has a distinct epistemic flavor. What looms large in Kenny’s scenario is the notion of skill. To reject the claim that Mary can, say, hit the top half of the board is not to deny that Mary might end up hitting that particular region of the target. What matters, instead, is that her hitting the top as opposed to the lower half of the board is a matter of chance—that she does not have enough control over her dart to favor one outcome over the other in the relevant way. The claim should not be that all cases in which distribution over disjunction fails are exactly like that: I can pick a raffle ticket with ease, but I cannot just pick a winning or a losing ticket, and it seems a bit strained (though perhaps not impossible) to talk about the presence and absence of a certain skill in this context. What the cases have in common, however, is that a certain path of action fails to be choosable since the agent does not really know how to implement it. Even if the agent does succeed at hitting the top part of the board or at picking a winning ticket, we scratch it up to luck since the action did not flow from an understanding of how the relevant outcome was to be achieved.

The view, in brief, is that the *can* of ability is essentially an existential quantifier over a set of available actions, and that an action is available to an agent just in case he or she is deemed to have sufficient understanding of how to achieve the relevant outcome. I have not said what it takes to possess such an understanding, and I am not going to start now. But it makes good sense to say that the relevant understanding is often intimately tied to the possession of certain skills, and that ordinary speakers are inclined to consider someone as skilled at doing something insofar as they believe the agent to have a good chance at succeeding in performing the relevant action, should he or she try to do it. All of this is reminiscent of the conditional analysis of abilities (see Mandelkern et al. 2017 for recent discussion). What matters for current purposes is the thought that even if this is so, this does not make the *can* in “Mary can hit the board” a conditional expression, any more than the correctness of a conditional analysis of dispositions would make the *might* in “The sugar might fully dissolve in this cup of water” express a conditional operator. Beliefs about abilities and dispositions—and thus, let us assume, conditional thoughts—certainly play a role for how we evaluate ordinary claims involving *can* and *might*, but they do so, I suggest, in a pragmatic way: by constraining the modal base that *can* and *might* take as input.

The proposal for ability *can* to be developed here effectively combines two layers of quantification: what is required is the existence of a proposition that entails the prejacent.

⁸The notion of refraining has been extensively studied in the stit-tradition (see in particular Belnap 1991 for a seminal detailed discussion); it is conceptually related to the notion of action negation, which has been extensively discussed in the PDL-tradition, starting with Meyer 1987a.

Appealing to two layers of quantification is certainly not an unfamiliar path—it is, in one way or another, the one pursued by Brown (1988), Hackl (1998), Horty (2001), and Portner (2009)—but what I am about to say here differs from previous stories in scope and matters of detail, not least because it makes sense of the previous observations about compulsion modals and free choice. One important point is that we can establish uniformity between possibility modals. Epistemic and deontic possibility modals, we said, require the existence of certain state of affairs, and we can think of a state of affairs as a maximal consistent proposition. As such, all modals quantify over sets of propositions—but they differ when it comes to negation and distribution over disjunction because they quantify over *different kinds* of propositions (more details to follow).

And there is another major issue: any story of the kind told here must account for the fact that the result of negating ability *can* in discourse seems to be surprisingly strong. To wit, the following remark about Mary’s upcoming throw sounds strange (as Mandelkern et al. (2017) observe):

- (13) Mary cannot hit the bullseye. ?? But she might.

This is a bit surprising if the first sentence in (13) were to simply negate the availability of a certain action, since the fact that no available action entails that Mary hits the bullseye leaves the possibility of her hitting the bullseye by accident open. Negating *can*, in brief, seems to say that the prejacent is incompatible with every choosable action, and not just that its negation is compatible with every choosable action.

This is not the first time that a plausible analysis of a certain empirical phenomenon seems to make negation too weak. Semantic approaches to the free choice effect need to explain why disjunction behaves classically under negation:

- (14) Mary may not hit the top or the lower half of the board.
 a. \rightsquigarrow Mary may not hit the top half of the board.
 b. \rightsquigarrow Mary may not hit the lower half of the board.
- (15) Mary cannot hit the top or the lower half of the board.
 a. \rightsquigarrow Mary cannot hit the top half of the board.
 b. \rightsquigarrow Mary cannot hit the lower half of the board.

The obvious moral to draw here would be that “Mary may hit the top or the lower half of the board” cannot *mean* “Mary may hit the top half of the board, and she may hit its lower half” since this would make the negation of a disjunctive possibility way too weak, as demonstrated by (14) and (15). This conclusion, however, is arguably not irresistible if we adopt a “bilateral” semantic system that distinguishes between, roughly, two distinct foundational semantic concepts: truth-making and false-making for instance, or support and rejection, or (as I will do here) positive and negative updating.⁹ We will later see that such an approach also captures the fact that negated *can* seems to entail *won’t*.

The present proposal explains Kenny’s observation almost by design, since of course one may know how to hit the board without knowing how to hit any particular region of the board—no wonder that distribution over disjunction fails if the *can* is the *can* of ability. Our setup also allows us to expect that action descriptions will interact in interesting ways with logical operators such as negation and disjunction. The first difference

⁹See for instance Aher 2012 and Willer 2015, 2017, 2018. The observation that negation poses problems for semantic accounts of free choice goes back at least to Alonso-Ovalle 2006 and Klinedinst 2007.

that matters here is the one between being able to bring about a disjunctive state of affairs and being able to bring about the state of affairs described by the first disjunct, or to bring about the one described by the second disjunct. In the latter case we have a disjunction taking scope over two action descriptions, and I will show why embedding such constructions under *can* gives rise to the free choice inference.

The second important difference is between being unable not to fail to do something and being unable not to refrain from doing something. The former will turn out to be equivalent to a negated ability attribution. But it is the latter which captures what we have in mind when we use compulsion modals: the circumstances compel the subject to actively *not* to do something. This difference, in turns out, corresponds to the one between negating that an agent does something and affirming that the agent actively does not do something.

The upshot is that our observations about ability *can* can be explained in terms of how logical operators such as disjunction and negation scopally interact with action descriptions. The underlying ideas can be elaborated in a variety of ways but to fix ideas I opt for a dynamic implementation.¹⁰ This is first and foremost because earlier I promised an account of why the *can* of ability sometimes not only distributes over disjunction but also licenses the even stronger free choice inference, and I know how to tell such a story in a dynamic setting using simple scope distinctions. If others can see a path toward implementing the spirit of the upcoming proposal in their preferred static setting—may it be truth-conditional, inquisitive, or what have you—all the power to them.

What I am about to say is closest in spirit to Update Semantics. Veltman (1996) proposes to treat semantic values as update functions on information carriers (sets of possible worlds). An update function for a simple modal propositional language, with “ \diamond ” interpreted as epistemic *might*, looks as follows:

- (1) $s \uparrow p = \{w \in s : w \in \llbracket p \rrbracket\}$
- (2) $s \uparrow \neg\phi = s \setminus (s \uparrow \phi)$
- (3) $s \uparrow (\phi \wedge \psi) = (s \uparrow \phi) \cap (s \uparrow \psi)$
- (4) $s \uparrow \diamond\phi = \{w \in s : s \uparrow \phi \neq \emptyset\}$

Of special importance here is the idea that atomic sentences are designed to add the proposition expressed to the input state, while the modal *might* effectively tests whether the prejacent is compatible with the input state.

The framework to be developed here differs from Update Semantics in various important respects while preserving its treatment of modals as tests on a modal domain. First, I will treat terminal semantic values as update *relations* (rather than update *functions*) to account for the logic of disjunction and also distinguish between a positive (acceptance inducing) and a negative (rejection inducing) update relation to account for the logic of negation. Second, I will introduce modal selection functions *e*, *d*, *a*, and so on, to distinguish between different modal flavors. Their range uniformly consists of propositions, but not all propositions are created equal, and it is the special kind of propositions that ability *can* quantifies over that will explain why it sometimes resists distribution over disjunction and plays with negation in a very special way. Let us go through the details.

¹⁰The classical sources of inspiration: Discourse Representation Theory (Kamp 1981; Kamp and Reyle 1993; Kamp et al. 2011), Dynamic Predicate Logic (Groenendijk and Stokhof 1991a), File Change Semantics (Heim 1982), Update Semantics (Veltman 1985, 1996).

3 Details

I will first provide the formal details of the proposal (Section 3.1) and then discuss its implications (Section 3.2). Section 3.3 explores what the framework has to say about raising modal verbs other than ability *can*, with a focus on the distinction between “evaluative” and “deliberative” interpretations of deontic modals.

3.1 Framework

Terminal semantic values, I said, are positive or negative update relations between states understood as sets of possible worlds and if D is the domain of individuals and A is a basic predicate of arity n , then each possible world is a (potentially partial) function from $D^n \times A$ to the classical truth-values. Predicate expressions denote a mapping from objects to positive or negative update relations between states. Here it will do no harm to focus on unary predicate expressions and to set aside context and tense. The semantic entry for a predicate such as *run* is stated below. A note on the notation: to avoid confusion, I will continue to use “[.]” to label the more familiar semantic objects such as propositions and predicate extensions; in contrast, “[.]” maps expressions to *dynamic* semantic values such as positive and negative update relations or—in the case of predicates—functions from objects to such relations.

$$(16) \quad \begin{array}{l} \text{a. } [run](x)(+) = \lambda s \lambda t. t = \{w \in s : w(x, run) = 1\} \\ \text{b. } [run](x)(-) = \lambda s \lambda t. t = \{w \in s : w(x, run) = 0\} \end{array}$$

For a state t to be positively related to another state s via some unary basic predicate β , given some object x , is for it to include all those worlds in s at which x is in the positive extension of β . And for a state to be thus negatively related is for it to include all those worlds w in s at which x is in the anti-extension of β . Note here that we allow in principle for an object to be neither in the extension nor in the anti-extension of a predicate at some world, that is, it does not trivially follow from the fact that $w(x, \beta) \neq 1$ that $w(x, \beta) = 0$.

Possible worlds directly fix the extension and anti-extension of basic predicates but we also need to consider complex predicates such as “not hitting the board” or “hitting the top or the lower half of the board.” As a first step we will define the extension of a predicate β , $[[\beta]]$, at some possible world w as follows:

$$(17) \quad [[\beta]](w) = \lambda x. \langle \{w\}, \{w\} \rangle \in [\beta](x)(+)$$

The thought here is that a sentence is true at a possible world just in case it is informationally redundant with respect to the perfect information about that world. The extension of a predicate β at a world is then the set of objects to which the predicate truthfully applies at that world. A familiar caveat: in an update-based dynamic setting like the one developed here, modal expressions have content but not truth-conditional content (see van Benthem 1986 and Groenendijk and Stokhof 1991b for classical discussion) and so we restrict the domain of $[[\cdot]]$ to the non-modal fragment of our target language.

Given some object, a predicate then more generally selects each world from the input state at which the predicate truthfully applies (if it has positive flavor) or not (if it has negative flavor). One wrinkle: it matters whether we talk about, say, the property

of bringing about some disjunctive state of affairs $s_1 \vee s_2$ or about the disjunction of bringing about s_1 and of bringing about s_2 . So whenever we are concerned with a complex predicate we will use “ $|\cdot|$ ” as a notational device for marking scopal relations and say the following (to avoid clutter, I will omit scope marking whenever we are concerned with a basic predicate):

$$(18) \quad \begin{array}{l} \text{a. } \llbracket \beta \rrbracket(x)(+) = \lambda s \lambda t. t = \{w \in s : x \in \llbracket \beta \rrbracket(w)\} \\ \text{b. } \llbracket \beta \rrbracket(x)(-) = \lambda s \lambda t. t = \{w \in s : x \notin \llbracket \beta \rrbracket(w)\} \end{array}$$

We can thus distinguish between “ $|\beta \vee \gamma|$ ” and “ $|\beta| \vee |\gamma|$ ” as well as between “ $|\neg\beta|$ ” and “ $\neg|\beta|$.” These scopal distinctions will turn out to be crucial to account for our previous observations about disjunction and negation. To indicate this for the case of negation: feeding some state s into $\llbracket \neg\beta \rrbracket(x)(+)$ will identify those worlds w in s at which x is in $\llbracket \neg\beta \rrbracket$, that is, the anti-extension of β ; in contrast, $\llbracket \neg|\beta| \rrbracket(x)(+)$ would identify those worlds w in s at which x fails to be in $\llbracket \beta \rrbracket$. These outputs are not trivially identical if we allow for partial worlds. Let me make this more precise by stating the semantics of the connectives. (Note that since $\llbracket \cdot \rrbracket$ is restricted to the non-modal fragment of our target language, so must be $|\cdot|$.)

Sentential negation (\neg) denotes a function $\#$ mapping any bilateral update relation $f: \{+, -\} \mapsto S \times S$, with S being the set of sets of states, to another according to the following rule:

$$(19) \quad \begin{array}{l} \text{a. } \#f(+) = \lambda s \lambda t. \langle s, t \rangle \in f(-) \\ \text{b. } \#f(-) = \lambda s \lambda t. \langle s, t \rangle \in f(+) \end{array}$$

This effectively says that if ϕ is a sentence, then s is positively related to t via “ $\neg\phi$ ” just in case s is negatively related to t via ϕ ; and s is negatively related to t via “ $\neg\phi$ ” just in case s is positively related to t via ϕ .

Given two bilateral update relations f and g , the functions denoted by conjunction (\wedge) and disjunction (\vee) are defined as follows:

$$(20) \quad \begin{array}{l} \text{a. } (f \sqcap g)(+) = \lambda s \lambda t. \exists u : \langle s, u \rangle \in f(+) \text{ and } \langle u, t \rangle \in g(+) \\ \text{b. } (f \sqcup g)(+) = \lambda s \lambda t. \langle s, t \rangle \in f(+) \text{ or } \langle s, t \rangle \in g(+) \\ \text{c. } (f \sqcap g)(-) = \lambda s \lambda t. \langle s, t \rangle \in f(-) \text{ or } \langle s, t \rangle \in g(-) \\ \text{d. } (f \sqcup g)(-) = \lambda s \lambda t. \exists u : \langle s, u \rangle \in f(-) \text{ and } \langle u, t \rangle \in g(-) \end{array}$$

For a state s to be positively related to t via “ $\phi \wedge \psi$ ”, there must be a state u to which s is positively related via the first conjunct, and u in turn must be positively related to t via the second conjunct. A state is positively related to another via a disjunction just in case they are positively related via one of the disjuncts. The negative entries enforce the validity of DeMorgan’s Laws.

Following standard protocol (see Rooth and Partee 1983) we can extend the rules for sentential connectives so that they apply beyond the basic conjoinable type of bilateral update relations. Say that if b is a conjoinable type, so is $\langle a, b \rangle$ for every a . And if f and g are derived conjoinable types in $D_{\langle a, b \rangle}$, then the functions denoted by negation, conjunction, and disjunction work as follows:

$$(21) \quad \text{(i) } \#f = \lambda x. \#f(x), \text{ where } x \text{ is of type } a$$

- (ii) $f \sqcap g = \lambda x. f(x) \sqcap g(x)$, where x is of type a
- (iii) $f \sqcup g = \lambda x. f(x) \sqcup g(x)$, where x is of type a

Suppose, for instance, that f and g are properties. Then $f \sqcap g$ amounts to the property of being both f and g , and it maps each object x to one of two transition instructions: the positive one, which adds to s the information that x has both properties; and the negative one, which adds to s the information that x lacks at least one of these properties. Relatedly, $\#f$ amounts to the property of not being f , and it maps any object x to one of two transition instructions: the positive one, which adds to s the information that x has the property of not being f ; and the negative one, which adds to s the information that x lacks the property of not being f . Let me now explain the proposal for modal expressions.

Modal selection functions map each state to a set of states, which then serves as the modal domain. I shall follow Portner (2009) in assuming that modal verbs are sometimes raising and sometimes control verbs, and that volitional modals, including the *can* of ability, belong to the latter category. But let us start with the familiar conception of modal expressions as sentential operators. In its positive flavor, a possibility modal uniformly requires that its prejacent be compatible with the modal domain; in its negative flavor, it uniformly requires that the prejacent be incompatible with the modal domain. One minor wrinkle: updates are defined over sets of worlds, not over sets of sets of worlds. We take care of this issue by universally quantifying over what I shall call the “specifications” of a modal domain (recall that a choice function γ on some collection X of nonempty sets assigns to each set x in that collection some element $\gamma(x)$ of x):

A set of worlds u is a *specification* of a set of sets of worlds X , $u \triangleright X$, just in case there is a choice function γ on X such that $u = \{z : \exists x \in X. \gamma(x) = z\}$.

In other words, a specification of a modal domain is simply the result of selecting exactly one possible world from each proposition in that modal domain.

The basic analysis then is that a positive update with a possibility modal tests whether each specification of the modal domain is compatible with the prejacent. A negative update with a possibility modal tests whether each specification of the modal domain is incompatible with the prejacent. Precisely, and letting “ $\underline{_}$ ” represent any non-empty state:

- (22) a. $[\diamond_f \phi](+) = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright f(s). \langle u, \emptyset \rangle \notin [\phi](+)\}$
- b. $[\diamond_f \phi](-) = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright f(s). \langle u, \underline{_} \rangle \notin [\phi](+)\}$

Passed tests return the original input state; failed tests return the empty set.

It is straightforward to extend the basic proposal so that it covers cases in which modals act as control predicates and thus take their subjects as arguments: we generalize the rule for “ \diamond ” and “ \square ” in the same way as we did for negation in (21). I shall also follow Portner (2009) in assuming that in such cases the modal selection function takes the subject as an extra argument.¹¹ Thus we get:

¹¹To make this more precise, observe that the sentential operator “ \diamond ” effectively denotes a function $[\diamond]$ that takes a modal selection function $m : S \mapsto \mathcal{P}(S)$ and a bilateral update relation $r : \{+, -\} \mapsto S \times S$ as input and returns another bilateral update relation. Say that a modal selection function is *suitable* for a bilateral update relation just in case it is a mapping from states to sets of states, and that a derived modal selection function is suitable for a conjoinable type in $D_{\langle a, b \rangle}$ just in case it is a function from D_a to modal selection functions suitable for objects in D_b . Then if f is a conjoinable type in $D_{\langle a, b \rangle}$ and m is suitable for f , $[\diamond]_m f = \lambda x. [\diamond]_{m(x)} f(x)$, where x is of type a , as expected. (23) is nothing but the instantiation of this schema where $D_{\langle a, b \rangle}$ is the set of properties.

- (23) a. $[\Diamond_f \beta](x)(+) = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright f(x, s). \langle u, \emptyset \rangle \notin [\beta](x)(+)\}$
 b. $[\Diamond_f \beta](x)(-) = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright f(x, s). \langle u, \underline{_} \rangle \notin [\beta](x)(+)\}$

Specifically, I shall assume that the selection function for the *can* of ability maps a subject x and state s to a set of actions: the actions that are available to x given the circumstances described by s .

Earlier I said that it would be desirable to continue treating possibility and necessity modals as duals so that “ $\Box_f \beta$ ” =_{df} “ $\neg \Diamond_f \neg \beta$ ” for all β . Here are the semantic entries for the necessity operator that are derived from (23):

- (24) a. $[\Box_f \beta](x)(+) = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright f(x, s). \langle u, \underline{_} \rangle \notin [\beta](x)(-)\}$
 b. $[\Box_f \beta](x)(-) = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright f(x, s). \langle u, \emptyset \rangle \notin [\beta](x)(-)\}$

Positive updates with necessity modals require that the modal domain not be negatively related to a consistent proposition via the prejacent. Negative updates with necessity modals require that the modal domain not be negatively related to the inconsistent proposition via the prejacent.

This is all that is needed. Let me finish things off by stating what it takes to update a state with a formula and the notion of entailment:

- (25) Let s be a state, $\phi_1, \dots, \phi_n, \psi$ be sentences. Define:
- The result of *updating* s with ϕ is defined as $s \uparrow \phi = \bigcup \{t : \langle s, t \rangle \in [\phi](+)\}$
 - s *supports* ϕ , $s \models \phi$, just in case $s \uparrow \phi = s$
 - ϕ_1, \dots, ϕ_n *entails* ψ , $\phi_1, \dots, \phi_n \models \psi$, iff for all s , $s \uparrow \phi_1, \dots, \phi_n \models \psi$

The result of updating s with ϕ is just the union of the states positively related to s via ϕ . A state supports ϕ just in case the result of updating s with ϕ returns s . A sequence of sentences entails ψ just in case the result of updating any s with that sequence supports ψ .

3.2 Output

It is straightforward to see how distribution over disjunction failures can arise for ability *can*: an action such as hitting the board may be deemed available to some agent without any more specific action being deemed available to that agent, and that is of course what seems to be going on in Kenny’s scenario. To get this into clearer view, take the following minimal list of possible worlds:

- w_t : Mary hits the top part of the board.
- w_l : Mary hits the lower part of the board.
- $w_{\bar{h}}$: Mary misses the board.

If s_c is the contextually provided common ground and a the selection function for ability *can*, then it makes good sense to say that $\{w_t, w_l\} \in a(\text{Mary}, s_c)$ but also $\{w_t\} \notin a(\text{Mary}, s_c)$ and $\{w_l\} \notin a(\text{Mary}, s_c)$. It follows that every specification of $a(\text{Mary}, s_c)$ is compatible with Mary’s hitting the board; but not every such specification is compatible with Mary’s hitting the top half of the board, and not every such specification is compatible with Mary’s hitting the lower half of the board.

The framework thus predicts that a state of information may support the claim that an individual can bring about a disjunctive state of affairs without being able to bring about any of its disjuncts. At the same time, it has the resources needed to explain why disjunctions under ability *can* sometimes license the free choice effect. What matters here is the contrast between the following two modal predications (assume here and throughout that β and γ are atomic predicate expressions):

$$\begin{aligned}
(26) \quad & \langle s, t \rangle \in [\Diamond_a |\beta \vee \gamma|](x)(+) & & \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \emptyset \rangle \notin [|\beta \vee \gamma|](x)(+)\} & & \text{(i)} \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \exists w \in u. x \in [|\beta \vee \gamma|](w)\} & & \text{(ii)} \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \exists w \in u. \langle \{w\}, \{w\} \rangle \in [|\beta \vee \gamma|](x)(+)\} & & \text{(iii)} \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \exists w \in u. \langle \{w\}, \{w\} \rangle \in [|\beta|](x)(+) \sqcup [|\gamma|](x)(+)\} & & \text{(iv)} \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \exists w \in u. x \in w(\beta) \text{ \underline{or} } x \in w(\gamma)\} & & \text{(v)} \\
(27) \quad & \langle s, t \rangle \in [\Diamond_a |\beta| \vee |\gamma|](x)(+) & & \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \emptyset \rangle \notin [|\beta| \vee |\gamma|](x)(+)\} & & \text{(i)} \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \emptyset \rangle \notin [|\beta|](x)(+) \sqcup [|\gamma|](x)(+)\} & & \text{(ii)} \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \emptyset \rangle \notin [|\beta|](x)(+) \text{ \underline{and} } \langle u, \emptyset \rangle \notin [|\gamma|](x)(+)\} & & \text{(iii)} \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \exists w \in u. x \in [|\beta|](w) \text{ \underline{and} } \exists w \in u. x \in [|\gamma|](w)\} & & \text{(iv)} \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \exists w \in u. x \in w(\beta) \text{ \underline{and} } \exists w \in u. x \in w(\gamma)\} & & \text{(v)}
\end{aligned}$$

(26) ascribes to x the ability to bring about a state that is either a β - or a γ -state. This requires the existence of an action α in $a(x, s)$ which entails that x is in the extension of $\beta \vee \gamma$, which in turn requires that $x \in w(\beta)$ or $x \in w(\gamma)$ for all $w \in \alpha$. Since the existence of such an action does not entail the existence of β - or of a γ -entailing action, the *can* in (26) fails to distribute over disjunction.

In (27), in contrast, the disjunction takes wide scope over two action descriptions—what is said here is that the subject can bring about the one state of affairs, or the other state of affairs—and this in fact does entail that both actions are available to the agent. For suppose that there is no action entailing that β applies to x : then there is a specification u of $a(x, s)$ that includes no world at which x is in the extension of β and thus $\langle u, \emptyset \rangle \in [|\beta|](x)(+)$. But then $\langle u, \emptyset \rangle \in [|\beta| \vee |\gamma|](x)(+)$ —recall that two states are positively related via a disjunction just in case they are thus related via one of the disjuncts—which violates the test conditions of (27). For parallel reasons, the test conditions of (27) require that the subject can bring about a γ -state (cf. (iii)). Together with the relational analysis of disjunction, our semantics for the possibility modal thus predicts the free choice effect; using simple scope distinctions, we can explain why the modal sometimes fails to distribute over disjunction.¹²

¹²It is fair to ask—as an anonymous reviewer does—why overt disjunction seems to block a Kenny reading. Consider:

(†) Mary can hit the top or bottom of the board.

It is hard to see how one could describe Mary's dart skills like this if one did not believe that Mary is skilled enough to hit the top half of the board, or skilled enough to hit the lower half of the board. Indeed, out of the blue a free choice reading of overt disjunction seems strongly preferred (see also Nouwen forthcoming, footnote 3). A plausible explanation would be that overt predicate disjunction has a strong tendency to take wide scope over $|\cdot|$ and is thus naturally interpreted as coordinating two action descriptions rather than as describing some brought about disjunctive state of affairs. As we have seen, such a wide scope reading licenses the free choice inference. On this view, Kenny's observation is real, but it actually matters that the ability attribution at play involves a covert rather than an overt disjunction operator.

Our explanation of the free choice effect generalizes to other modals, since nothing we said in the previous paragraph depends on the nature of the modal domain. The current framework also accounts for the observation that disjunctive possibilities behave classically under the scope of negation. Start by observing that negated *cans* require the incompatibility of the prejacent with any available action:

$$\begin{aligned}
(28) \quad & \langle s, t \rangle \in [\Diamond_a \beta](x)(-) \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \underline{\quad} \rangle \notin [\beta](x)(+)\} & (i) \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \forall w \in u. x \notin w(\beta)\} & (ii) \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \forall w \in u. w(x, \beta) \neq 1\} & (iii)
\end{aligned}$$

Assuming that a state of affairs is possible only if it is at least compatible with an agent's actions, we thus explain why *can't* entails *won't*. Moreover, our semantics for disjunction predicts that negating a disjunctive possibility is to negate the possibility of both disjuncts:¹³

$$\begin{aligned}
(29) \quad & \langle s, t \rangle \in [\Diamond_a |\beta| \vee |\gamma|](x)(-) \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \underline{\quad} \rangle \notin [|\beta| \vee |\gamma|](x)(+)\} & (i) \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \underline{\quad} \rangle \notin [\beta](x)(+) \sqcup [\gamma](x)(+)\} & (ii) \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \underline{\quad} \rangle \notin [\beta](x)(+) \text{ and } \langle u, \underline{\quad} \rangle \notin [\gamma](x)(+)\} & (iii) \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \forall w \in u. x \notin [\beta](w) \text{ and } \forall w \in u. x \notin [\gamma](w)\} & (iv) \\
& \text{iff } t = \{w \in s : \forall u \triangleright a(x, s). \forall w \in u. x \notin w(\beta) \text{ and } \forall w \in u. x \notin w(\gamma)\} & (v)
\end{aligned}$$

The key observation here: if β were compatible with some action in $a(x, s)$, then $\langle u, \underline{\quad} \rangle \in [\beta](x)(+)$ and thus—again given the semantics for disjunction— $\langle u, \underline{\quad} \rangle \in [|\beta| \vee |\gamma|](x)(+)$ for some $u \triangleright a(x, s)$, which means that (29) is rejected by s (assuming that s is consistent). The same would hold if γ were compatible with some action in $a(x, s)$. I thus conclude that our bilateral setup gets the basic facts about negated *cans* straight. Let me now explore their relation with compulsion modals in a bit more detail.

Earlier I claimed that the negation of an ability modal does not obviously entail the corresponding compulsion statement. Start by recalling the update relation denoted by a negated ability attribution:

$$\begin{aligned}
(30) \quad & [\Diamond_a \beta](x)(-) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \underline{\quad} \rangle \notin [\beta](x)(+)\} & (i) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \forall w \in u. x \notin w(\beta)\} & (ii) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \forall w \in u. w(x, \beta) \neq 1\} & (iii)
\end{aligned}$$

On this analysis, to say that Mary cannot hit the bullseye is to say that she fails to hit the bullseye no matter what she does. Compare this with what is communicated by saying that Mary cannot but not hit the bullseye, as analyzed in (31):

(This explanation, I should add, is compatible with the familiar possibility of an ignorance reading that is made salient by continuing the above example with, for instance, “but I don't know which” (Kamp 1978): in this case, the disjunction is sentential and takes wide scope over two ability attributions, that is, the sentence is of the form $\ulcorner \Diamond_a \phi \vee \Diamond_a \psi \urcorner$.)

¹³We could also let the disjunction take narrow scope here, but this does not lead to an interesting difference. It is easy to verify that for any consistent s , $\langle s, s \rangle \in [\Diamond_a |\beta| \vee |\gamma|](x)(-)$ just in case no $\alpha \in a(x, s)$ includes a world at which $x \in w(\beta)$, or a world at which $x \in w(\gamma)$.

$$\begin{aligned}
(31) \quad & [\Box_a | \neg \beta](x)(+) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \underline{\quad} \rangle \notin [|\neg \beta](x)(-)\} & (i) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \forall w \in u. x \in [|\neg \beta](w)\} & (ii) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \forall w \in u. \langle \{w\}, \{w\} \rangle \in [|\neg \beta](x)(+)\} & (iii) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \forall w \in u. \langle \{w\}, \{w\} \rangle \in \#[\beta](x)(+)\} & (iv) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \forall w \in u. \langle \{w\}, \{w\} \rangle \in [\beta](x)(-)\} & (v) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \forall w \in u. w(x, \beta) = 0\} & (vi)
\end{aligned}$$

Analyzed like this, “Mary cannot but not hit the bullseye” requires that Mary refrains from hitting the bullseye no matter what she does. And refraining from doing β , recall, is stronger than failing to do β : the former requires the presence of an intention causing behavior incompatible with doing β , while the former is perfectly compatible with the presence of an intention to do β . In other words, being unable to hit the bullseye is compatible with seriously trying, while having the compulsion not to hit the bullseye is not.

On the story told here, a compulsion to do β does correspond to an inability: the inability not to *refrain* from doing β . This is compatible with “Mary cannot but not hit the bullseye” being stronger than “Mary can’t hit the bullseye.” The latter articulates an inability as well, but the inability not to *fail* to hit the bullseye. In fact, we can distinguish the salient reading of “Mary cannot but not hit the bullseye” from one that has a less compulsive flavor along the lines of “Mary cannot but fail to hit the bullseye.” If we capture the latter as being of the form in (32), then it turns out to be equivalent to a simple negated *can*:

$$\begin{aligned}
(32) \quad & [\Box_a \neg |\beta](x)(+) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \underline{\quad} \rangle \notin [|\neg \beta](x)(-)\} & (i) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \underline{\quad} \rangle \notin \#[\beta](x)(-)\} & (ii) \\
& = \lambda s \lambda t. t = \{w \in s : \forall u \triangleright a(x, s). \langle u, \underline{\quad} \rangle \notin [\beta](x)(+)\} & (iii) \\
& = [\Diamond_a \beta](x)(-) & (iv)
\end{aligned}$$

As such, we once again can appeal to scope differences to account for some subtle intuitions about ability *can*. And I submit that this is not merely a technical result but tracks a real distinction between different kinds of omissions.

Finally, let me observe that distribution over disjunction as well as duality failures disappear if the modal domain consists of state descriptions understood in the following sense:

A state description is a singleton set $\{w\}$ such that for all $x \in D$ and $\beta \in A$,
 $w(x, \beta) = 1$ or $w(x, \beta) = 0$.

This is easy to see since the difference between a negated possibility and the necessity of the corresponding negation entirely depends on the gap between an object’s failing to be in the extension of a predicate and its falling into the anti-extension of that predicate. Relatedly, distribution over disjunction failures depend on there being a proposition entailing a disjunction without entailing at least one of its disjuncts, which falls away if we are only dealing with singleton sets. If we now assume that epistemic and deontic modals quantify over state rather than action descriptions, it follows straightaway that they contrast with ability *can* in the ways we observed (and regardless of whether or not they are control or raising predicates).

The framework proposed here offers an explanation of a range of puzzling observations about the *can* of ability while preserving a distinct sense of uniformity across modals: across the board, *can* requires possibility understood as compatibility with some modal domain. What makes the *can* of ability so special is that it requires the possibility of an action and I have demonstrated that this account helps us understand why the *can* of ability is so special.

3.3 Raising Beyond *Can*

I have treated the *can* of ability as a control verb, and I note that everything I have said here lives happily with the claim that other modal expressions may be control verbs as well. This is important, not least since one often hears it in the literature on deontic modality that we are well-advised to recognize an “ought to do” that is irreducible to the “ought to be” familiar from standard deontic logic.¹⁴ Schroeder (2011), for instance, details that there is an evaluative, raising reading of *ought* on which it says that, if things were ideal, some proposition would be the case (the “ought to be” reading); but there is also a deliberative, control reading of *ought* on which it relates an agent to some property (the “ought to do” reading).¹⁵ Relatedly, Portner (2009) suggests that some deontic modals are raising and other control verbs, and Thomason (1981) argues that there must be, in addition to impersonal uses of *ought*, a personal *ought* which relates a property and an individual. Let me briefly demonstrate that the framework makes perfect sense of these suggestions (Section 3.3.1) and compare the resulting picture with some prominent alternative analyses of the “ought to do” (Section 3.3.2).

3.3.1 Ought and Control: Analysis

In the current framework the positive entry for the raising reading of *ought* is captured by (33a) and the one for the control reading by (33b):

- (33) a. $[\Box_d\phi](+) = \lambda s\lambda t.t = \{w \in s : \forall u \triangleright d(s). \langle u, \underline{\lambda} \rangle \notin [\phi](-)\}$
 b. $[\Box_d\beta](x)(+) = \lambda s\lambda t.t = \{w \in s : \forall u \triangleright d(x, s). \langle u, \underline{\lambda} \rangle \notin [\beta](x)(-)\}$

In (33a), the modal has the type of a sentential operator; in (33b) it has the type of a predicate.

While the *ought*-readings in (33a) and (33b) have a lot in common, they differ enough to make sense of the kind of data that have traditionally motivated the view that deliberative *oughts* differ from their evaluative cousins (see, among others, Geach 1982; Horty and Belnap 1995; Schroeder 2011). Consider:

- (34) a. Bob ought to dance with Mary.

¹⁴See, e.g., the overview by Hilpinen and McNamara (2012) for a recent discussion of standard deontic logic, who also note that the seminal contribution by von Wright (1951) actually does not treat deontic operators as sentential operators (as one does in standard deontic logic) but rather as applying to expressions for action types.

¹⁵The opposing view that all deliberative *oughts* are reducible to evaluative *oughts* is voiced by, among many others, Meinong (1917), Chisholm (1964), and Williams (1991); see also Chrisman (2012) and Finlay and Snedegar (2014), who critically discuss Schroeder’s (2011) arguments. The other extreme—that all evaluative *oughts* reduce to deliberative *oughts*—is less popular in the literature, but see Geach 1982 and also D’Altan et al. 1996. Wedgwood (2006, 2007) and Broome (2013) distinguish between an evaluative and a deliberative sense of *ought* but hold that the latter relates agents to a proposition.

- b. Mary ought to dance with Bob.

On the evaluative reading given by (33a), (34a) and (34b) are equivalent: the former effectively requires that Bob dance with Mary at all deontically ideal worlds; since the relation of dancing is symmetric, this is to require that Mary dance with Bob at all deontically ideal worlds, which is just what (34b) demands. And yet there is a sense on which (34a) and (34b) fail to be equivalent: it is easy to imagine social conventions, for instance, that require Bob to dance with Mary, but which do not impose any such obligation on Mary. This sense is captured by the control reading of *ought* in (33b). Note, in particular, that the deontic selection function in (33b) takes the subject as an additional argument, and it may very well be that, for some input state s , $d(\text{Bob}, s) \neq d(\text{Mary}, s)$ —Bob’s obligations need not be Mary’s obligations—and so s may support (34a) without supporting (34b), or vice versa, if the *ought* at play behaves as in (33b).

And nothing prevents us from shaping our subject-sensitive deontic selection function so that it accounts for other intuitions that have been articulated in the literature. Harman (1986) holds that the deliberative *ought*, but not its evaluative cousin, is “agent-implicating” in the sense that it implies the agent’s ability to actually implement the action described by the prejacent. To impose this requirement, we would say that for all states s , individuals x , and predicates p if for all $u \in d(x, s)$, $x \in \llbracket \beta \rrbracket(w)$ for all $w \in u$, then for some $v \in a(x, s)$, $x \in \llbracket \beta \rrbracket(w)$ for all $w \in v$ —in other words, we restrict the subject-sensitive deontic selection function so that personal *oughts* imply personal abilities. Relatedly, we may follow Horty (2001) and assume that the actions available to an agent can be ranked in terms of their utility, and take this as the starting point of how we think of what an agent ought to do by letting $d(x, s)$ be sensitive to the utility of the actions that are available to an agent (rather than, say, to deontically ideal histories, which may be all that evaluative *oughts* care about). Specifically, let us follow Kratzer (1981, 1991, 2012) and say that the deontically ideal states are determined by a contextually provided ordering source, which is a set of propositions that induces an ordering on the states in the modal domain: $d(x, s)$ is then just the set of state descriptions that are ideal in light of this ranking. We may then say that p is in the ordering for $d(x, s)$ just in case for all actions that x can perform (according to s) and that do not result in p there is a higher ranked action that does result in p , and all actions that are ranked even higher also result in p .¹⁶ Whether any of these claims about deliberative *oughts* is correct is a question that must be left to another day: the point that matters here is that the framework presented is rich enough to accommodate not only the intuition that there is a substantial difference between deliberative and evaluative *oughts* but also a variety of specific proposals about the logic and semantics of what an agent ought to do, deliberatively speaking.

It is worth noting here that we can distinguish between an evaluative and a deliberative *ought* and still maintain that all deontic modals quantify over state rather than action descriptions—in fact, this seems just right since both kinds of *oughts* happily distribute over disjunction and play with negation and permission in a classical fashion. Obliga-

¹⁶This proposal for linking deontic obligations with the utility of available actions is Horty’s (2001) reaction to the so-called “gambling problem” for the proposal that all deliberative *oughts* reduce to evaluative *oughts* (which Horty labels the “Meinong-Chisholm thesis”). I am effectively suggesting here that broadly Kratzerian approaches to deontic modal semantics can implement at least certain key ideas from the deontic logic literature. See Horty 2014 for some reflections on the scope as well as the limits of pursuing this strategy when it comes to implementing ideas from approaches to deontic logic that are based on default logic.

tions exist, on this view, because certain states count as deontically ideal. The difference between the evaluative and the deliberative *ought* arises because what counts as a deontically ideal state, deliberatively speaking, depends on the subject, but not so if the *ought* at play has an evaluative flavor, and in general we are free to impose constraints on subject-sensitive modal selection functions that play no role for their subject-insensitive cousin. In brief, the “ought to do” and the “ought to be” differ in how their modal domains are determined in discourse and reasoning, not in the kind of entities that constitute those modal domains.

3.3.2 Ought and Control: Comparisons

The resulting account of *ought* (and *may*) is one moment that distinguishes the dynamic framework developed here from prominent existing analyses that draw their distinctly dynamic flavor from their rooting in the Propositional Dynamic Logic (PDL) tradition.¹⁷ PDL is a formal system for reasoning about the input/output behavior of programs of computers. In addition to regular propositional formulas, PDL features $[\alpha]\phi$ and $\langle\alpha\rangle\phi$, where the former says that no matter how program α is executed, ϕ will be the case afterwards (alternatively, that α *leads* to ϕ) and the latter says that there is some way to execute program α such that ϕ will be the case afterwards (alternatively, that α *may* lead to ϕ). Programs are binary relations: we say that a pair $\langle s_1, s_2 \rangle$ is in the interpretation of a program α just in case the execution of α can cause the transition from s_1 to s_2 . If we now think of input/output states as sets of possible worlds and of actions as programs, actions turn out to be relations between sets of possible worlds.

On the basis of atomic actions we may construct several *complex actions*, for example $\sim\alpha$ (not- α), $\alpha\sqcup\beta$ (α or β), $\alpha\&\beta$ (α and β), and $\alpha;\beta$ (α followed by β). While the question of how action negation is to be interpreted is a matter of some controversy (see e.g. Broersen 2004), it is common to interpret the remaining complex action operators as union, intersection and relative product, respectively.

Based on a variant of PDL, Meyer (1987a, 1987b) proposes a dynamic deontic logic that introduces a propositional constant V —standing for *violation*—which is said to be true in a state just in case a violation occurs in that state. He then proposes the following reduction of the notion of an action α being forbidden (inspired by Anderson 1967; see also Kanger 1971):

$$(R1) F\alpha \leftrightarrow [\alpha]V$$

An action α is forbidden ($F\alpha$) just in case every execution of α leads to a violation. Together with the strong interdefinability of prohibition and permission—and action is permitted just in case it is not forbidden—we arrive (together with some basic principles of dynamic logic) at the following reduction of permission:

$$(R2) P\alpha \leftrightarrow \langle\alpha\rangle\neg V$$

It is permitted to do α just in case there is a way to execute α that leads to a state in which no violation occurs.

One immediate problem with this proposal is that there is no easy way to generalize it so that it yields a uniform analysis of all existential modal notions, since it is not clear

¹⁷See, e.g., Harel 1984 for a classical introduction to dynamic logic.

what should count as violation when it comes to the areas of knowledge and of ability. But even setting aside this fundamental concern, Meyer’s (1987a, 1987b) reduction fails to deal with a puzzle that the framework proposed in this paper handles with grace: the issue of free choice permission. Recall examples such as the following:

- (35) a. Mary may/is permitted to hit the top or the lower half of the board.
 b. Mary may/is permitted to hit the top half of the board.
 c. Mary may/is permitted to hit the lower half of the board.

The familiar observation is that (35a) implies the conjunction of (35b) and (35c). This is not predicted by classical deontic logic since the possibility of a disjunction is perfectly compatible with the impossibility of one of its disjuncts. And it is not predicted by (R2) either, since the fact that, if presented with a choice between α and β , I may steer clear of trouble is perfectly compatible with, say, the choice of α inevitably resulting in a violating state. In contrast, free choice effects follow immediately in the dynamic framework developed here; I give both the positive and the negative update rule for the “may/is permitted to do” (everything I will say here also applies to the evaluative interpretation of permission):

- (36) a. $[\Diamond_d\beta](x)(+) = \lambda s\lambda t. t = \{w \in s : \forall u \triangleright d(x, s). \langle u, \emptyset \rangle \notin [\beta](x)(+)\}$
 b. $[\Diamond_d\beta](x)(-) = \lambda s\lambda t. t = \{w \in s : \forall u \triangleright d(x, s). \langle u, \perp \rangle \notin [\beta](x)(+)\}$

We already saw in Section 3.2 that given (36a), (35a) entails the conjunction of (35b) and (35c) under the reasonable assumption that the disjunction in (35a) takes wide scope over two action descriptions: roughly, given arbitrary β , if there were no permissible state entailing that β applies to x , then there would be a specification u of $d(x, s)$ including no world at which x is in the extension of β and thus $\langle u, \emptyset \rangle \in [\beta](x)(+)$, and so $\langle u, \emptyset \rangle \in [|\beta| \vee |\gamma|](x)(+)$, violating the test conditions of (35a) if analyzed as predicating “ $\Diamond_d|\beta| \vee |\gamma|$ ” of Mary.

Broersen (2004) proposes a different reduction:

$$(R3) P\alpha \leftrightarrow [\alpha] - V_P$$

Here V_P represents a special kind of violation (“lack of explicit permission”) and there is no assumption of strong interdefinability. Intuitively, the idea is that an action α is permissible just in case all ways of executing α are alright.¹⁸ This reduction does deliver free choice, since “ $P(\alpha \cup \beta)$ ” now says that every way of executing $\alpha \cup \beta$ is alright, and both doing α (in some way or other) and doing β (in some way or other) are ways of executing $\alpha \cup \beta$; so every way of executing α and every way of executing β are alright, and hence we have $P(\alpha \cup \beta) \rightarrow P\alpha \wedge P\beta$. For the same reason, however, Broersen’s (2004) reduction implies

$$P\alpha \rightarrow P(\alpha \& \beta)$$

since doing $\alpha \& \beta$ is a way doing α and so if any execution of the latter is alright, so must be any execution of the former. This is certainly a highly controversial result about permission (or deontic *may*), and it is deeply problematic if we are trying to make sense of the free choice phenomenon. It follows from everything we said so far:

¹⁸Meyer (1987a) also considers a reduction in this spirit, though he adds the condition that there exists a permissible execution of α . For related proposals, see, e.g., Segerberg 1982 and van der Meyden 1996.

$$P(\alpha \cup \beta) \rightarrow P(\alpha \& \beta)$$

And the consensus in the literature on free choice is that the following is perfectly consistent:

- (37) You may have a soup or a salad, but you may not have both.

A permission to do α or β implies a conjunction of permissions, but not the permission of a conjunction.¹⁹ The conclusion must be that Broersen’s (2004) reduction does not get the basic facts about free choice permission right.

Another issue left unaddressed by (R3) is the issue of “double prohibition.” Not only does disjunction behave non-classically under the scope of deontic possibility operators, it also behaves classically again if the possibility is negated:

- (38) a. Mary may not/is not permitted to hit the top or the lower half of the board.
 b. Mary may not/is not permitted to hit the top half of the board.
 c. Mary may not/is not permitted to hit the lower half of the board.

(38a) implies both (38b) and (38c). This is problematic if we are aiming for a reduction that makes free choice a matter of entailment (as does R3) and assume a classical notion of negation: $\neg P(\alpha \cup \beta)$ is now much too weak, true if either α or β may lead to a violation.²⁰ In contrast, the bilateral framework proposed here makes the right predictions. Roughly, if there were a permissible state entailing that β applies to x , then every specification u of $d(x, s)$ would include a world at which x is in the extension of β and thus $\langle u, \underline{x} \rangle \in [\beta](x)(+)$, and so $\langle u, \underline{x} \rangle \in [|\beta| \vee |\gamma|](x)(+)$, violating the negative test conditions of (38a).

Nothing said here demonstrates that we cannot come up with a PDL-inspired analysis of deontic permission that delivers all the goodies. It does, however, show that not all theories with a dynamic flavor are created equal, and that the dynamic framework proposed here makes non-trivial contributions to the literature on deliberative and evaluative deontic modality. More generally, a dynamic treatment of action terms as denoting relations between states can only be one building block among many of a successful modal analysis. I will say more about this issue when I return to ability *can* in Section 4.1.

4 Loose Ends

This section briefly touches upon a few remaining issues. I begin by contextualizing the story told here by exploring some salient alternative strategies (Section 4.1) and then answer some resulting questions about the path pursued here (Section 4.2). Section 4.3 notes some issues that need to be left to another day for a comprehensive discussion.

¹⁹The problematic inference may be blocked by suitably restricting what counts as a permissible execution of an action type, for instance by adding some operator that selects the “minimal” executions of an action: α and β , if performed individually, are minimal executions of $\alpha \cup \beta$ in the relevant sense, while $\alpha \& \beta$ is not (see Dignum et al. 1996). No such maneuvers are necessary to make sense of free choice effects in the framework developed here.

²⁰I note that the problem of negation also affects recent sophisticated approaches to free choice permission, and regardless of whether they are in the dynamic logic tradition (see e.g. Ju and van Eijck 2019) or not (see e.g. Anglberger et al. 2016).

4.1 Alternatives

I have said that ability modals quantify over actions, and I have made certain assumptions about actions in order to account for some of the data that interest us here: that the availability of some action does not guarantee the availability of a more specific action—in response to Kenny’s observation about distribution over disjunction failures—and that there is a distinction between failing to do something and refraining from doing something—in response to our observations about duality failures. These are not unfamiliar claims and in fact the basic analysis of ability and compulsion is highly sympathetic to the stit-semantic analysis of these notions by Horty (2001). But it would be a mistake to think that all the puzzles I have discussed here can just be solved by getting the logic of actions straight—in fact, if what is said is right here we need just as much reflection about modals, negation, disjunction, and scope to sort things out. To make the point stick, let me begin by briefly reviewing how far we can get with a basic stit-semantic analysis of ability *can* and how much further the framework proposed here is taking us.

Stit-semantics is cast against the background of a theory of indeterministic time (see Prior 1967 and Thomason 1970 for seminal contributions). Start with a frame $\langle Tree, < \rangle$ where *Tree* is a non-empty set of moments and $<$ is a strict partial ordering of these moments such that for all $m, m',$ and m'' from *Tree*, if $m' < m$ and $m'' < m$, then either $m' = m''$ or $m' < m''$ or $m'' < m'$ (no backward branching). A maximal set of linearly ordered moments from *Tree* is a *history* and a single moment may be an element of distinct histories: $H^m = \{h : m \in h\}$ are the histories “passing through m ” and when h belongs to H^m , we speak of a moment/history pair of the form m/h as an *index*. Atomic sentences are assigned truth-values relative to such indices and basic operators receive their meanings familiar from classical intentional semantics. For instance, we say (with \mathcal{M} being some model):

- (i) $\mathcal{M}, m/h \models p$ iff $v(m/h, p) = 1$
- (ii) $\mathcal{M}, m/h \models \neg\phi$ iff $\mathcal{M}, m/h \not\models \phi$
- (iii) $\mathcal{M}, m/h \models (\phi \wedge \psi)$ iff $\mathcal{M}, m/h \models \phi$ and $\mathcal{M}, m/h \models \psi$
- (iv) $\mathcal{M}, m/h \models (\phi \vee \psi)$ iff $\mathcal{M}, m/h \models \phi$ or $\mathcal{M}, m/h \models \psi$
- (v) $\mathcal{M}, m/h \models \diamond\phi$ iff $\mathcal{M}, m/h' \models \phi$ for some history $h' \in H^m$

The proposition expressed by ϕ at some moment m , $|\phi|_{\mathcal{M}}^m = \{h \in H^m : \mathcal{M}, m/h \models \phi\}$, is the set of histories from H^m in which that sentence is true.

The crucial innovation is the introduction of a *stit*-operator, capturing the idea that an agent α sees to it that ϕ means that the truth of ϕ is guaranteed by an action performed by α . Embellish a frame $\langle Tree, < \rangle$ with a set *Agent* of agents and a function *Choice* mapping each agent α and moment m into a partition $Choice_{\alpha}^m$ of the set H^m of histories through m . $Choice_{\alpha}^m(h)$ is the equivalence class from $Choice_{\alpha}^m$ that contains h . Each such partition represents a choice available to α at m , and we may then define a standard *stit*-operator as follows:

- (vi) $\mathcal{M}, m/h \models [\alpha \textit{ stit} : \phi]$ iff $Choice_{\alpha}^m(h) \subseteq |\phi|_{\mathcal{M}}^m$

Additional constraints can be imposed on the *Choice* function. These need not detain us here; what matters for current purposes is that we can now offer a very natural definition of ability: what an agent is able to do is identical with what it is possible for the agent

to choose. The claim that α can see to it that ϕ is thus to be analyzed as follows:

$$\diamond(\alpha \textit{ stit}: \phi)$$

The ability to bring about ϕ amounts to the existence of a choice that guarantees the truth of ϕ .²¹

The stit-semantic analysis and the story proposed here have something important in common: both essentially analyze ability *can* as combining existential with universal quantification. In the stit-framework, the universal force stems from the presence of the *stit*-operator. This delivers a solution to Kenny’s problem and it also predicts that negated abilities do not amount to compulsions (assuming that the latter are analyzed as a *stit* scoping under a necessity operator). But it does not *by itself* take care of all the data that the story told here addresses. For starters, we observed in Section 1 that one not only needs to explain why distribution over disjunction fails for ability *can*, but also why it does hold for deontic *may* and for epistemic *might*. This is unexpected in stit-semantics, since there is no reason to think that *stit*-operators cannot scope under deontic or epistemic possibility modals and so we would expect distribution over disjunction failures across the (modal) board.²² In contrast, the story proposed here does not only explain why ability *can* differs from its epistemic and deontic cousins in how it distributes over disjunction, but also offers just the kind of explanation that one would hope for: in terms of differences between modal domains.

Another point at which the framework goes beyond the basic stit-story—or, for that matter, any classical framework that analyzes ability *can* as combining existential with universal quantification—is the treatment of negation. The result of negating ability *can* in stit-semantics seems too weak: it denies the existence of a choice that guarantees the truth of prejacent, but not that the prejacent is possible at all. This, recall from Section 2, raises the question of why constructions such as (39a) have the air of a contradiction:

- (39) a. Al cannot hit the bullseye. ?? But it is possible that he does (by chance).

²¹This particular analysis of ability *can* is based on the “Chellas stit,” introduced by Horty and Belnap (1995), who draw their inspiration from Chellas (1969) and also consider other stit-flavored operators. Horty and Pacuit (2017) offer an analysis of ability *can* that is based on a *kstit*-operator and couched in a labeled stit-framework. The resulting innovations do not affect the upcoming discussion of how stit-semantics differs from the framework proposed here, though the *kstit*-based analysis is similar in spirit to the proposal made here insofar as both impose distinct epistemic constraints on what it takes for an action to be available to an agent.

²²A similar problem arises for the proposal by Mandelkern et al. (2017) to alleviate Kenny’s concern using the distinction between general and specific abilities. The idea here is that in Kenny’s scenario, we ascribe to Mary the *general* ability to hit the board: in a suitable proportion of normal situations, hitting the board is something that Mary can do. This can be true without there being a suitable proportion of normal situations in which hitting the top half of the board is something that Mary can do, and without there being a suitable proportion of normal situations in which hitting the lower half of the board is something that Mary can do. The presence of an unarticulated genericity operator would thus explain why the possibility modal fails to distribute over disjunction. But the question now is why a generic quantifier should not be present when we talk about what Mary might or is permitted to do with the board: there is, after all, no immediate reason to think that we cannot talk about what someone, in general, might or may do. The contrast between ability *can* and epistemic *might* and deontic *may* remains unexplained. Likewise for a supervaluationist response: Mary can definitely hit the board, but it is not definitely true that she can hit the top/the bottom half of the board (see again Mandelkern et al. 2017). While close in spirit to the story told here—or, for that matter, any traditional story that combines existential with universal generalization—the question remains essentially the same: why supervalue in interesting ways when we talk about ability *can*, but not when we talk about *might* and *may*?

- b. $\neg\Diamond[\alpha \textit{ stit} : p] \wedge \Diamond p$

(39a) receives its most natural formalization in (39b), which is perfectly consistent in stit-semantics. This leaves it a bit of a puzzler why (39a) sounds marked, not least because structurally identical constructions sound fine if we change the first conjunct’s modal flavor. Specifically, while it is strange to say that Al might end up doing something that he cannot do, there is nothing strange about saying that he might end up doing something that he is not permitted not do.

- (40) Al may not hit the bullseye. But it is possible that he does (by chance).

All of this suggests that there is something specifically about the meaning of ability *can* that makes its negation incompatible with possibility.

Mandelkern et al. (2017) go so far and take the data surrounding (39a) as a silver bullet against all existential-universal analyses of ability *can*. What I have demonstrated here is that these approaches are actually defensible if coupled with a suitably sophisticated account of negation—one that is couched in a bilateral semantics that distinguishes between positive and negative updating (or two other basic semantic concepts, such as acceptance and rejection). This allows us to say that a *positive* update with a possibility modal requires each specification of the modal domain to be compatible with the prejacent—a requirement that is essentially analogous to the stit-semantic analysis—while a *negative* update with a possibility modal tests whether each specification of the modal domain is incompatible with the prejacent. To my knowledge, this is the first response to the problem of negation on behalf of the existential-universal analysis of ability *can*.

Finally, we noted that every satisfying story about ability *can* not only needs to explain why it sometimes resists distribution over disjunction but also sometimes licenses—like other modals do—the even stronger free choice inference. Free choice inferences are invalid on the classical stit-analysis of ability *can*.

- (41) a. Mary can hit the top or the lower half of the board.
 b. $\Diamond[\alpha \textit{ stit} : p \vee q]$
 c. $\Diamond([\alpha \textit{ stit} : p] \vee [\alpha \textit{ stit} : q])$
 d. Mary can hit the top half of the board, and she can hit the lower half of the board.
 e. $\Diamond[\alpha \textit{ stit} : p] \wedge \Diamond[\alpha \textit{ stit} : q]$

The observation is that (41a) implies (41d), which is formalized in (41e). (41a) allows for the two analyses in (41b) and (41c), yet it is easy to verify that neither of them entails (41e). Of course, nothing stops us from trying to supplement stit-semantics with a pragmatic account of free choice. What matters here is that the framework developed in this paper readily accounts for the puzzling observation that ability modals sometimes resist distribution over disjunction and sometimes even license free choice inferences: free choice effects arise whenever a disjunction coordinates two action descriptions (roughly analogous to (41c)), whereas distribution over disjunction failures become a possibility whenever we are attributing the ability to bring about a disjunctive state of affairs (as we do in (41b)). The required scope distinctions, to be clear, can be drawn with the help of a stit-operator, but we still need a story about how possibility modals interact with

disjunctions to get the facts about free choice straight. The framework developed here offers such a story.

The same points could be made if we were to look at other existing analyses of ability in the literature, and even if we look beyond the ones already mentioned earlier (Brown 1988, Hackl 1998, Portner 2009).²³ For instance, Segerberg’s (1989, 1992) PDL-based analysis of ability—which adds a “bringing about” operator to action terms in PDL and then treats ability in a stit-style fashion as the possibility to reliably bring about the truth of some proposition—does not go beyond stit-semantics when it comes to the lack of an explanation of why ability *can* contrasts with its epistemic and deontic cousins, the trouble with negation, and the silence on free choice; nor does the KARO framework developed by van Linder et al. (1998) (see also van der Hoek et al. 1999 and Meyer et al. 1999), which analyzes ability as the knowledge of having a certain practical possibility, or Elgesem’s (1997) semantics for the “bringing-it-about-that,” which introduces a selection function f such that $f_i(w, X)$ is the set of those worlds where i realizes the ability that i has in w to bring about the goal of X , and then says that i is able to bring about X at w just in case $f_i(w, X) \neq \emptyset$. The general moral is that a plausible theory about action and agency can only be one component among many in a comprehensive story about ability *can*: every such story needs to talk about deontic and epistemic modals, and it needs to offer a sophisticated story about negations, disjunctions, and their sensitivity to scopal interactions. It is through developing and integrating all of these components into a coherent picture that the framework presented here makes a genuine and positive contribution to the literature.

None of this is meant to show that alternative frameworks for analyzing the *can* of ability ought to be abandoned. In fact, some of the key ideas developed here may be transferred into alternative settings to broaden their coverage. Return to the stit-framework, and consider the problem with negation: that negated *cans* are too weak because they only deny the existence of a choice guaranteeing the truth of prejacent, not the possibility of the prejacent. Structurally similar problems affect all existential-universal analyses of ability *can*, including (in principle) the present proposal, and in response I suggested a bilateral semantic system that distinguishes between, roughly, two distinct foundational semantic concepts: truth-making and false-making for instance, or support and rejection, or (the path chosen here) positive and negative updating. The idea can also be put to good use in a stit-framework, and the obvious choice would be to distinguish between a truth-making (\models^+) and a false-making (\models^-) relation between an index and a sentence (given some model). Begin with the obvious entries for atomic sentences and negation:

$$\begin{array}{ll} \text{(i-a)} & \mathcal{M}, m/h \models^+ p \text{ iff } v(m/h, p) = 1 \\ \text{(i-b)} & \mathcal{M}, m/h \models^- p \text{ iff } v(m/h, p) = 0 \end{array} \quad \begin{array}{ll} \text{(ii-a)} & \mathcal{M}, m/h \models^+ \neg\phi \text{ iff } \mathcal{M}, m/h \models^- \phi \\ \text{(ii-b)} & \mathcal{M}, m/h \models^- \neg\phi \text{ iff } \mathcal{M}, m/h \models^+ \phi \end{array}$$

Conjunction and disjunction may be given their straightforward positive and negative truth-conditions so that DeMorgan’s Laws remain in force. The first steps toward getting negated *cans* right is to couple the already familiar stit-treatment of the possibility operator with a separate negative clause:

²³The conditional analysis proposed by Mandelkern et al. (2017) is explicitly designed to make negated *cans* sufficiently strong, but in exchange collapses inability and compulsion.

- (v-a) $\mathcal{M}, m/h \models \diamond\phi$ iff $\mathcal{M}, m/h' \models^+ \phi$ for some history $h' \in H^m$
(v-b) $\mathcal{M}, m/h \models \diamond\phi$ iff $\mathcal{M}, m/h' \models^- \phi$ for all histories $h' \in H^m$

For a negated possibility to be true at a moment m (and history h), the negation of the prejacent must be true at all histories passing through m . Finally, we refine the *Choice*-operator as follows:

- (vi-a) $\mathcal{M}, m/h \models^+ [\alpha \textit{ stit}: \phi]$ iff $\textit{Choice}_\alpha^m(h) \subseteq |\phi|_{\mathcal{M}}^m$
(vi-b) $\mathcal{M}, m/h \models^- [\alpha \textit{ stit}: \phi]$ iff $\textit{Choice}_\alpha^m(h) \cap |\phi|_{\mathcal{M}}^m = \emptyset$

Failing to see to it that ϕ , on this proposal, is stronger than failing to guarantee the truth of ϕ : the former now requires that every choice available to the agent is incompatible with the truth of ϕ .

This solves the problem with negation: to say that α cannot see to it that p at m/h is to say that p is incompatible with every possible choice available to α at m (i.e. with each element of \textit{Choice}_α^m). Assuming again that this can only be so if p is actually impossible at m/h , we account for why cases such as (39a) are off. And we can preserve the crucial difference between failing to be able to see to it that p and the compulsion to see to it that not- p by once again assuming that there are truth-value gaps: the fact that $v(m/h, p) \neq 1$ does not trivially entail that $v(m/h, p) = 0$. On this picture, the inability to see to it that ϕ (at m/h) amounts to every available choice being incompatible with ϕ (meaning: for every $h' \in H^m$, for every $h'' \in \textit{Choice}_\alpha^m(h')$, ϕ fails to be true at m/h''); the compulsion not to ϕ at m/h , in contrast, amounts to every available choice entailing the proposition expressed by $\ulcorner \neg\phi \urcorner$ at m (meaning: for every $h' \in H^m$, $\textit{Choice}_\alpha^m(h') \subseteq |\neg\phi|_{\mathcal{M}}^m$).

The idea of tracing the contrast between ability *can* and its epistemic and deontic cousins to differences between modal domains is in principle applicable to stit-semantics as well. The simplest way to implement this thought—though other avenues are available—would be to add to our models modal selection functions mapping each index m/h to a set of indices $f(m/h)$. Modals would be existential or universal quantifiers over these sets of indices. Importantly, if we now insist that the *Choice*-function maps each element m/h of the epistemic and deontic modal domain to the singleton set $\{m/h\}$ (but to choices available at m/h when it comes to ability *can*), the presence of a *stit*-operator under the scope of an existential modal would block distribution of disjunction for ability *can* but not for epistemic *might* and deontic *may*. Explaining why the *Choice*-function should behave like this will require additional conceptual groundwork—more than can be done here—but there is no general reason to think that such an explanation could not be provided.

More resistant to a simple transfer of ideas is the treatment of free choice. While stit-frameworks are rich enough to reproduce the scopal distinctions that matter for the story told here—between cases in which a disjunction coordinates two action descriptions and those in which there is an action of bringing about a disjunctive state of affairs—we already noted that none of these actually entail the free choice reading in classical stit-semantics. The solution proposed here crucially relies on a *dynamic* semantics in which terminal semantic values are relations between information carriers, and in particular on the idea that (i) a disjunction relates an input state s_1 to another state s_2 just in case one of its disjuncts does and (ii) possibility modals encode a prohibition against the prejacent's relating the modal domain to the absurd state. It remains to be seen whether a similarly successful and elegant analysis of free choice effects, their interaction with

negation, and their dependency on scope can be developed in a non-dynamic setting such as classical stit-semantics.²⁴

Still, the framework proposed here faces its own explanatory challenges. The contrast between ability *can* and its epistemic and deontic cousins, I have claimed, seems best explained in terms of differences between their respective modal domains. But it is fair to ask why these differences between modal domains should exist. Let me make a few additional remarks about this issue.

4.2 Modal Domains

The claim that modal (and other intensional) expressions affect the interpretation of their prejacent is anything but unprecedented. Condoravdi (2002) argues that modals combine with untensed sentences and fix the temporal orientation of their prejacent directly; similar claims have been made about non-finite embedding attitudes (see e.g. Abusch 2004). Klecha (2016) has recently refined these ideas and concludes that constraints on temporal interpretation flow from the nature of the selected modal base. The idea that ability *can* gives its prejacent an agentive flavor since its modal domain consists of a set of available actions is just another spin on this idea.

I have not much to say in support of the claim that epistemic and deontic modals effectively quantify over ordinary possible worlds. This is the default position in the literature, and the fact that it fits our observations about how these modals distribute over disjunction and interact with negation indicates that the view is on the right track. (Recall also from Section 3.3 that this position is perfectly compatible with drawing a meaningful distinction between evaluative and deliberative interpretations of deontic modals.) If there is any source of concern, it should be that we arguably can draw attention to the possibility of Mary’s hitting the board without drawing attention to her hitting any particular part of the board. Likewise, we may give someone the explicit permission to hit the board without specifying which area that individual may hit. I will sketch a way to accommodate this intuition momentarily.

It is perhaps less obvious why certain modals should quantify over action descriptions in the way they are understood here. I suggest that we can see why this should be so if we consider the role that ability *can* and its dual play for our everyday rational activities. We frequently engage in means-end reasoning: what to do in order to achieve certain goals. Such reasoning, it is often said, carries a distinct practical flavor in that it is designed to result in an action. To do that, however, it is not enough to simply identify some doing that would, in theory, realize the goal under consideration: knowing that cutting the correct wire will disarm the bomb does not leave me with a plan about how to disarm the bomb if I do not know what the correct wire is; knowing that switching the timer off will do the trick as well might be more useful, assuming I know how to do that. As Ford (2016) puts it, for a doing to play the role of a means to an end in practical deliberation, it must be a *practical* possibility: an action that the deliberating agent can actually implement.

I take it to be uncontroversial that possibility modals are designed to highlight their prejacent as a distinct possibility in discourse and reasoning. I add that ability *can* is designed to highlight its prejacent as a practical possibility—a possibility that may figure

²⁴Avenues worth exploring here: combining our modified stit-framework with state-based (Aloni ms.) or with homogeneity-based (Goldstein forthcoming) approaches to free choice effects.

in the attributee’s practical deliberation, should it turn out to be conducive to his or her practical interests. And of course if such possibilities can only exist when the agent understands how to implement them, it is no surprise that the *can* of ability cannot quantify over maximally complete state descriptions but must allow for cases in which some less opinionated proposition depicts a possibility without some more opinionated proposition also depicting a possibility.²⁵

The appeal to means-end reasoning also seems to be illuminating when it comes to the use of gappy possible worlds in modeling the domain of ability modals. Start with the epistemic and deontic case: necessity modals here articulate certainties and obligations, respectively, and of course if p is impossible, then its negation is a certainty; and if p is impermissible, it is required that p does not occur. Compulsion modals, I propose, effectively assert that the attributee has a non-negotiable end: whatever else the agent is doing, he or she must find the means conducive to that end.²⁶ The key observation here: the fact that doing β is not an option—cannot figure as a means in practical deliberation—does not trivially make not doing β an end worth pursuing, for the plain reason that not doing β may be entirely irrelevant to one’s practical concerns. A practical impossibility, in other words, does not trivially amount to a practical necessity—the proposed distinction between failing to do β and not doing β is responsive to exactly this fact.

The underlying picture here is that language allows us to draw the distinctions that we want to make and that modal quantifier domains, in particular, are responsive to the kind of modality that the corresponding modal expressions are designed to track. My claim then is that the framework proposed here does not only account for the empirical facts on the ground. Modal thought and talk, I said, frequently pertains to what is practically possible or necessary since we often engage in thought and talk about means and ends, and we often do so saying (or thinking) that such-and-such *can* or *must* be done. The previous remarks on what makes a possibility practical—and on how practical impossibilities relate to practical necessities—suggest that the assumptions we made about the modal domain of ability modals are not only empirically adequate but also track something real and important about the kind of modality that these modals are designed to express.

²⁵The earlier voiced intuition that the *can* of ability requires sufficient control over some outcome thus derives from what it takes for a possibility to count as a practical possibility. I note here that the notion of control also plays a role elsewhere in the literature: Farkas (1992), for instance, proposes that the strength of “subject obviation effects” (Ruwet 1984) depends on whether the occurrence of the situation described by the complement is under the subject’s control. Hargreaves (2005) identifies a class of “control verbs” in Kathmandu Newari that receive a distinct egophoric marking by default. Both phenomena have received analyses that rely on ideas from the *de se* literature (see Schlenker (2005) and Coppock and Wechsler (2018), respectively), and it would be interesting to see whether the *de se* also plays a role in the realm of practical possibilities. I must leave this question for another day, not least because the exact status of the *de se* for semantic theorizing is not entirely uncontroversial (see for instance Cappelen and Dever 2013).

²⁶There is a puzzle here that I can only mention without fully resolving it. Earlier we saw that the negation of “Mary can hit the bullseye” entails that Mary will not hit the bullseye; the same seems to be true if we say that Mary cannot but not hit the bullseye. Yet agents sometimes fail to realize their goals, and in particular Mary may end up hitting the bullseye even if she actively tries not to: so why is acknowledging this possibility infelicitous? This puzzle is not unique to compulsion modals: it arises in similar forms for strong necessity modals (Ninan 2005; Silk 2015, forthcoming) and commands (Mandelkern 2018). The general phenomenon seems to be that in some discourse situations certain possibilities somehow count as ruled out despite lack of hard evidence. While one may come up with discourse constraints that accommodate this empirical fact, there remains the difficult question of how we could ever rationally rule out possibilities without evidence. Marušić (2015), who focusses on ethical and epistemological dimensions of acts of promising, offers a highly relevant and interesting discussion.

4.3 Expansions

Some remaining issues must be left for another day. I have set aside the phenomenon of *actuality entailments*: in languages that make a clear distinction between perfective and imperfective aspect, perfective sentences containing ability *can* entail that the event described by the prejacent actually took place (Bhatt 1999). Addressing this issue goes far beyond the scope of the current investigation, but I note that everything said here is, for instance, compatible with the assumptions about tense and aspect that do the explanatory lifting in Hacquard’s (2006) analysis of this puzzling phenomenon.

It is also worth noting explicitly that the story told here, while relying on “|·|” to draw essential scope distinctions, does not appeal to an agency operator (such as the *stit*-operator) in its analysis of ability *can*. Part of the moral here is that such an operator is inessential for resolving the empirical puzzles that interest us here—including the observation that ability *can*, but not its epistemic or deontic cousins, resists distribution over disjunction. Nonetheless, it must be admitted that the absence of an agency operator brings certain analytical limitations in its wake. Perhaps most importantly, right now the framework cannot offer a *stit*-style analysis of the notion of refraining; nor does it fully capture the analytical differences and similarities between, say, the ability to *prevent* someone from entering a building, the ability to *make* someone enter a building, and the ability to *remove* someone from a building, since it can only treat the predicates at play here as primitive; nor can it capture the variety of ways in which an agent—or a group of agents—may be responsible for an event. Exploring all of these issues in more detail is a major task that must be left to another day, but there is no principle reason why an agency operator (or a variety of agency operators) could not be added to the existing framework, and nothing said so far precludes weaving insights from the existing literature on these topics into the existing analysis.²⁷

So far I have also not said much about the fact that not everything that can be the grammatical subject of *can* is an agent. I briefly mentioned one particular example earlier (repeated in (42a)) but it is easy to generate other cases:

- (42) a. The dart can hit the board.
b. The car can go 200 miles per hour.
c. This boulder can crush you.

Since darts, cars, and boulders cannot perform actions, these cases call for an explanation.

The most reasonable response, it seems to me, is that here the *can* has an ordinary circumstantial reading with state descriptions as its domain of quantification. This proposal receives some support by the fact that the cases in (42) can be rephrased using expletives:

- (43) a. It can happen that the dart hits the board.
b. It can happen that the car goes 200 miles per hour.
c. It can happen that this boulder crashes you.

²⁷Belnap 1991 is a seminal discussion of the thesis that refraining amounts to seeing to it that one does not see to it. While the detailed study of agent interactions is of general interest, it has been extensively explored in the analyses of normative positions (see Sergot 2012 and references therein for discussion). See Sergot 2014 for a detailed discussion that highlights, inter alia, the need for distinguishing between different notions of responsibility.

A similar rephrasing seems unavailable if the *can* is the one of ability. (44a) and (44b), for instance, do not seem to say the same thing (see Mandelkern et al. 2017 and Portner 2009 for related observations):

- (44) a. Mary can hit the board.
b. It can happen that Mary hits the board.

How then can the absence of an ability reading of the examples in (42) be explained? The simple suggestion is that the selection function for the *can* of ability is restricted: it only takes genuine agents as arguments. Granted, the notion of an agent that is at play here must then be more restrictive than the traditional, grammatical one of Agent external arguments. Specifically, the existing literature suggests that legitimate Agent external arguments need not per se be animate but rather must possess an internal setup suitable to generate the event or process denoted by their predicate on their own. Folli and Harley (2008), for instance, distinguish between Agent and Cause theta-roles but also argue on syntactic grounds that the DP “the train” in “The train whistled” is an Agent external argument despite being inanimate (see also Copley and Wolff 2014). But as Pietroski (1998) points out, it may very well be that speakers can draw distinctions that grammar cannot see, and I suggest that the difference between genuine agents and things that merely participate in an event is one of them.²⁸

Earlier I mentioned the intuition that one may draw attention to a possibility without bringing into view more specific possibilities. Telling you that Mary might be in France is not to draw attention to any particular region where she might be, and I may explicitly permit Mary to go to France without bothering to tell her where exactly she may go, and so on. It is fair to ask whether this fact is a problem for the story told here, since it suggests a potentially unwelcome parallel between ability talk and its epistemic and deontic cousins: in all of these discourse types can a coarse-grained possibility be contextually salient without any of its refinements being salient in the same way.

In response, we may accommodate the phenomenon under consideration here by noting that modal selection functions are underdetermined in discourse. What possibilities are to be treated as salient is not fully settled, and one thing that possibility modals do is to constrain the context in a distinct way: they sharpen the relevant selection function so that the prejacent is treated as compatible with the modal domain. I suggest that a proposition p counts as a contextually salient epistemic possibility just in case each candidate for the epistemic selection function e treats p as compatible with the epistemic modal domain. And of course a proposition can be salient—or be brought into salience—in this sense without the same being true of any more specific proposition.

I thus propose to treat the phenomenon of bringing possibilities into view as an instance of the more general process of contextual refinement, and it then makes sense to go a supervaluationist path and say that an inference is “super-valid” just in case it is valid on every contextual specification. Given everything we said about how the selection function for *might* and *may* differs from the one for *can*, it follows straightaway that distribution over disjunction is super-valid for epistemic and deontic modals, but not for

²⁸Given the criteria laid out here, even bona fide agents sometimes merely participate in events: “Mary can hit the board by accident,” for instance, can be rephrased as “It can happen that Mary hits the board by accident,” and it seems to follow that Mary can hit the top part of the board by accident, or that she can hit the lower half of the board by accident. Thanks to an anonymous reviewer for pointing out this possibility to me.

ability *can*. The fact that the contextual specifications need not agree on which disjunct is to be supported would then make the existence of these validities compatible with the considerations about salience that were articulated in the previous two paragraphs. Since the details are straightforward, let me move on and wrap up the discussion.

5 Conclusion

Ability modals differ from their epistemic and deontic cousins in interesting ways. The key message of this paper is that this fact can be explained in a uniform semantics for modals. Some of the central moves I made are certainly non-classical: terminal semantic values are bilateral update relations, and modal quantifier domains are sets of propositions. But this departure from how things are normally done allows us to account for the data using well-known explanatory tools and techniques: differences between modal inferential patterns correspond to differences between the underlying modal domains of quantification; scope distinctions account for subtle observations about distribution over disjunction failures, free choice inferences, and the interaction between *can*, negation, and compulsion modals. The non-classical machinery, in brief, ultimately brings a range of puzzling observations into familiar territory. Since the story told here is also conceptually sound—its key moves flow from reasonable assumptions about how practical possibility relates to practical necessity and how failing relates to refraining—I submit that what has been said here deserves being taken seriously.

References

- Abusch, Dorit. 2004. “On the Temporal Composition of Infinitives.” In *The Syntax of Time*, ed. Jacqueline Guéron and Jacqueline Lecarme, 27–53. Cambridge, MA: MIT Press.
- Aher, Martin. 2012. “Free Choice in Deontic Inquisitive Semantics (DIS).” *Lecture Notes in Computer Science* 7218: 22–31. http://dx.doi.org/10.1007/978-3-642-31482-7_3.
- Aloni, Maria. 2007. “Free Choice, Modals, and Imperatives.” *Natural Language Semantics* 15(1): 65–94. <http://dx.doi.org/10.1007/s11050-007-9010-2>.
- . ms. “FC Disjunction in State-based Semantics.” Manuscript, University of Amsterdam.
- Alonso-Ovalle, Luis. 2006. *Disjunction in Alternative Semantics*. Ph.D. thesis, University of Massachusetts, Amherst.
- Anderson, Alan Ross. 1967. “Some Nasty Problems in the Formal Logic of Ethics.” *Noûs* 1(4): 345–360. <http://dx.doi.org/10.2307/2214623>.
- Anglberger, Albert J.J., Federico L.G. Faroldi, and Johannes Korbmacher. 2016. “An Exact Truthmaker Semantics for Permission and Obligation.” In *Proceedings of DEON 2016*, ed. Olivier Roy, Allard Tamminga, and Malte Willer, 16–31. London: College Publications.

- Barker, Chris. 2010. “Free Choice Permission as Resource-sensitive Reasoning.” *Semantics and Pragmatics* 3(10): 1–38. <http://dx.doi.org/10.3765/sp.3.10>.
- Belnap, Nuel. 1991. “Before Refraining: Concepts for Agency.” *Erkenntnis* 34(2): 137–169. <http://dx.doi.org/10.1007/BF00385718>.
- Belnap, Nuel, Michael Perloff, and Ming Xu. 2001. *Facing the Future: Agents and Choices in Our Indeterminist World*. New York: Oxford University Press.
- van Benthem, Johan. 1986. *Essays in Logical Semantics*. Dordrecht: Reidel. <http://dx.doi.org/10.1007/978-94-009-4540-1>.
- Bhatt, Rajesh. 1999. *Covert Modality in Non-Finite Contexts*. Ph.D. thesis, University of Philadelphia, Philadelphia.
- Broersen, Jan. 2004. “Action Negation and Alternative Reduction for Dynamic Deontic Logics.” *Journal of Applied Logic* 2(1): 153–168. <http://dx.doi.org/10.1016/j.jal.2004.01.007>.
- Broome, John. 2013. *Rationality Through Reasoning*. Malden, MA: Wiley Blackwell. <http://dx.doi.org/10.1002/9781118609088>.
- Brown, Mark A. 1988. “On the Logic of Ability.” *Journal of Philosophical Logic* 17(1): 1–26. <http://dx.doi.org/10.1007/bf00249673>.
- Cappelen, Herman, and Josh Dever. 2013. *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person*. New York: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199686742.001.0001>.
- Chellas, Brian F. 1969. *The Logical Form of Imperatives*. Ph.D. thesis, Stanford University.
- Chisholm, Roderick M. 1964. “The Ethics of Requirement.” *American Philosophical Quarterly* 1(2): 147–153. <http://www.jstor.org/stable/20009130>.
- Chrisman, Matthew. 2012. “‘Ought’ and Control.” *Australasian Journal of Philosophy* 90(3): 433–451. <http://dx.doi.org/10.1080/00048402.2011.611151>.
- Clarke, Randolph. 2014. *Omissions: Agency, Metaphysics, and Responsibility*. New York: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199347520.001.0001>.
- Condoravdi, Cleo. 2002. “Temporal Interpretation of Modals: Modals for the Present and for the Past.” In *The Construction of Meaning*, ed. David I. Beaver, Stefan Kaufmann, Brady Clark, and Luis Casillas, 59–88. Stanford, CA: CSLI Publications.
- Copley, Bridged, and Phillip Wolff. 2014. “Theories of Causation Should Inform Linguistic Theory and Vice Versa.” In *Causation in Grammatical Structures*, ed. Bridged Copley and Fabienne Martin, 11–57. Oxford: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199672073.003.0002>.

- Coppock, Elizabeth, and Stephen Wechsler. 2018. “The Proper Treatment of Egophoricity in Kathmandu Newari.” In *Expressing the Self: Cultural Diversity and Cognitive Universals*, ed. Minyao Huang and Kasia M. Jaszczolt, 40–57. Oxford: Oxford University Press. <http://dx.doi.org/10.1093/oso/9780198786658.003.0003>.
- Cross, Charles B. 1986. “‘Can’ and the Logic of Ability.” *Philosophical Studies* 50(1): 53–64. <http://dx.doi.org/10.1007/BF00355160>.
- D’Altan, Piero, John-Jules Ch. Meyer, and Roelf Johannes Wieringa. 1996. “An Integrated Framework for Ought-to-be and Ought-to-do Constraints.” *Artificial Intelligence and Law* 4(2): 77–111. <http://dx.doi.org/10.1007/BF00116787>.
- Dignum, Frank, John-Jules Ch. Meyer, and Roelf Johannes Wieringa. 1996. “Free Choice and Contextually Permitted Actions.” *Studia Logica: An International Journal for Symbolic Logic* 57(1): 193–220. <http://dx.doi.org/10.1007/BF00370675>.
- Elgesem, Dag. 1997. “The Modal Logic of Agency.” *Nordic Journal of Philosophical Logic* 2(2): 1–46.
- Farkas, Donka. 1992. “On Obviation.” In *Lexical Matters*, ed. Ivan Andrew Sag and Anna Szabolcsi, 85–109. Stanford, CA: CSLI Publications.
- Finlay, Stephen, and Justin Snedegar. 2014. “One Ought Too Many.” *Philosophy and Phenomenological Research* 89(1): 102–124. <http://dx.doi.org/10.1111/j.1933-1592.2012.00646.x>.
- Folli, Raffaella, and Heidi Harley. 2008. “Teleology and Animacy in External Arguments.” *Lingua* 118(2): 190–202. <http://dx.doi.org/10.1016/j.lingua.2007.02.004>.
- Ford, Anton. 2016. “On What Is in Front of Your Nose.” *Philosophical Topics* 44(1): 141–161. <http://dx.doi.org/10.5840/philtopics20164419>.
- Fox, Danny. 2007. “Free Choice and the Theory of Scalar Implicatures.” In *Presupposition and Implicature in Compositional Semantics*, ed. Uli Sauerland and Penka Stateva, 71–120. Hampshire: Palgrave Macmillan. http://dx.doi.org/10.1057/9780230210752_4.
- Franke, Michael. 2011. “Quantity Implicatures, Exhaustive Interpretation, and Rational Conversation.” *Semantics and Pragmatics* 4(1): 1–82. <http://dx.doi.org/10.3765/sp.4.1>.
- Fusco, Melissa. 2015a. *Deontic Disjunction*. Ph.D. thesis, UC Berkeley.
- . 2015b. “Deontic Modality and the Semantics of Choice.” *Philosophers’ Imprint* 15(28): 1–27. <http://dx.doi.org/2027/spo.3521354.0015.028>.
- Geach, Peter T. 1982. “Whatever Happened to Deontic Logic?” *Philosophia* 11(1-2): 1–12. <http://dx.doi.org/10.1007/BF02378802>.
- Geurts, Bart. 2005. “Entertaining Alternatives: Disjunctions as Modals.” *Natural Language Semantics* 13(4): 383–410. <http://dx.doi.org/10.1007/s11050-005-2052-4>.

- Giannakidou, Anastasia. 2001. “The Meaning of Free Choice.” *Linguistics and Philosophy* 24(6): 659735. <http://dx.doi.org/10.1023/A:1012758115458>.
- Giannakidou, Anastasia, and Eleni Staraki. 2012. “Ability, Action, and Causation: From Pure Ability to Force.” In *Genericity*, ed. Alda Mari, Claire Beyssade, and Fabio Del Prete, 250–275. Oxford: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199691807.003.0009>.
- Goldstein, Simon. Forthcoming. “Free Choice and Homogeneity.” *Semantics and Pragmatics*. <http://dx.doi.org/10.3765/sp.12.23>.
- Groenendijk, Jeroen, and Martin Stokhof. 1991a. “Dynamic Predicate Logic.” *Linguistics and Philosophy* 14(1): 39–100. <http://dx.doi.org/10.1007/BF00628304>.
- . 1991b. “Two Theories of Dynamic Semantics.” In *Lecture Notes in Artificial Intelligence: Logics in AI, European Workshop JELIA '90*, ed. Jan van Eijk, 55–64. Berlin: Springer-Verlag. <http://dx.doi.org/10.1007/BFb0018433>.
- Hackl, Martin. 1998. “On the Semantics of “Ability Attributions”.” <http://web.mit.edu/hackl/www/papers/files/AbilityAttributions.pdf>. Manuscript, MIT.
- Hacquard, Valentine. 2006. *Aspects of Modality*. Ph.D. thesis, MIT, Cambridge, MA.
- Harel, David. 1984. “Dynamic Logic.” In *Handbook of Philosophical Logic II: Extensions of Classical Logic*, ed. Dov Gabbay and Florian Guentner, 497–604. Dordrecht: Reidel. http://dx.doi.org/10.1007/978-94-009-6259-0_10.
- Hargreaves, David. 2005. “Agency and Intentional Action in Kathmandu Newar.” *Himalayan Linguistics* 5: 1–48. <https://escholarship.org/uc/item/36093022>.
- Harman, Gilbert. 1986. *Change in View*. Cambridge, MA: MIT Press.
- Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts, Amherst.
- Hilpinen, Risto, and Paul McNamara. 2012. “Deontic Logic: A Historical Survey and Introduction.” In *Handbook of Deontic Logic and Normative Systems*, ed. Dov Gabbay, John Horty, Xavier Parent, Ron van der Meyden, and Leon van der Torre, 3–136. London: College Publications.
- van der Hoek, Wiebe, Bernd van Linder, and John-Jules Ch. Meyer. 1999. “An Integrated Modal Approach to Rational Agents.” In *Foundations of Rational Agency*, ed. Michael J. Woolridge and Anand S. Rao, 133167. Dordrecht: Kluwer. http://dx.doi.org/10.1007/978-94-015-9204-8_7.
- Horty, John F. 2001. *Agency and Deontic Logic*. New York: Oxford University Press. <http://dx.doi.org/10.1093/0195134613.001.0001>.
- . 2014. “Deontic Modals: Why Abandon the Classical Semantics?” *Pacific Philosophical Quarterly* 4(95): 424–460. <http://dx.doi.org/10.1111/papq.12044>.

- Horty, John F., and Nuel Belnap. 1995. “The Deliberative Stit: A Study of Action, Omission, Ability, and Obligation.” *Journal of Philosophical Logic* 24(6): 583–644. <http://dx.doi.org/10.1007/BF01306968>.
- Horty, John F., and Eric Pacuit. 2017. “Action Types in STIT Semantics.” *The Review of Symbolic Logic* 4(10): 617–637. <http://dx.doi.org/10.1017/S1755020317000016>.
- Ju, Fengkui, and Jan van Eijck. 2019. “A Temporal Dynamic Deontic Logic.” *Journal of Logic and Computation* 29(2): 265284. <http://dx.doi.org/10.1093/logcom/exy008>.
- Kamp, Hans. 1973. “Free Choice Permission.” *Proceedings of the Aristotelian Society* 74: 57–74. <http://www.jstor.org/stable/4544849>.
- . 1978. “Semantics versus Pragmatics.” In *Formal Semantics and Pragmatics for Natural Languages*, ed. Franz Guenther and Siegfried Josef Schmidt, 255–287. Dordrecht: Reidel. http://dx.doi.org/10.1007/978-94-009-9775-2_9.
- . 1981. “A Theory of Truth and Representation.” In *Formal Methods in the Study of Language, Part I*, ed. Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, 277–320. Amsterdam: Mathematisch Centrum.
- Kamp, Hans, Josef van Genabith, and Uwe Reyle. 2011. “Discourse Representation Theory.” In *Handbook of Philosophical Logic*, 2nd ed., vol. 15, ed. Dov M. Gabbay and Franz Guenther, 125–394. Dordrecht: Springer. http://dx.doi.org/10.1007/978-94-007-0485-5_3.
- Kamp, Hans, and Uwe Reyle. 1993. *From Discourse to the Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. Dordrecht: Kluwer Academic Press. <http://dx.doi.org/10.1007/978-94-017-1616-1>.
- Kanger, Stig. 1971. “New Foundations for Ethical Theory.” In *Deontic Logic: Introductory and Systematic Readings*, ed. Risto Hilpinen, 36–58. Dordrecht: Reidel. http://dx.doi.org/10.1007/978-94-010-3146-2_2.
- Kenny, Anthony. 1975. *Will, Freedom and Power*. Oxford: Blackwell.
- . 1976. “Human Abilities and Dynamic Modalities.” In *Essays on Explanation and Understanding*, ed. Juha Manninen and Raimo Tuomela, 209–232. Dordrecht: Reidel. http://dx.doi.org/10.1007/978-94-010-1823-4_11.
- Klecha, Peter. 2016. “Modality and Embedded Temporal Operators.” *Semantics and Pragmatics* 9(9): 1–55. <http://dx.doi.org/10.3765/sp.9.9>.
- Klinedinst, Nathan. 2007. *Plurality and Possibility*. Ph.D. thesis, UCLA.
- Kratzer, Angelika. 1977. “What ‘Must’ and ‘Can’ Must and Can Mean.” *Linguistics and Philosophy* 1(3): 337–355. <http://dx.doi.org/10.1007/BF00353453>.
- . 1981. “The Notional Category of Modality.” In *Words, Worlds, and Contexts. New Approaches to Word Semantics*, ed. Hans-Jürgen Eikmeyer and Hannes Rieser, 38–74. Berlin: Walter de Gruyter & Co.

- . 1991. “Modality.” In *Semantics: An International Handbook of Contemporary Research*, ed. Armin von Stechow and Dieter Wunderlich, 639–650. Berlin: Walter de Gruyter & Co.
- . 2012. *Modals and Conditionals*. New York: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199234684.001.0001>.
- Kratzer, Angelika, and Junko Shimoyama. 2002. “Indeterminate Pronouns: The View from Japanese.” In *Proceedings of the Third Tokyo Conference on Psycholinguistics*, ed. Yukio Otsu, 1–25. Hituzi Sybobo.
- van Linder, Bernd, Wiebe van der Hoek, and John-Jules Ch. Meyer. 1998. “Formalising Abilities and Opportunities of Agents.” *Fundamenta Informaticae* 34(1,2): 53–101. <http://dx.doi.org/10.3233/FI-1998-341203>.
- Maier, John. 2013. “The Agentive Modalities.” *Philosophy and Phenomenological Research* 90(1): 113–134. <http://dx.doi.org/10.1111/phpr.12038>.
- Mandelkern, Matthew. 2018. “You Must! Maybe You Won’t.” In *Proceedings of Deontic Logic and Normative Systems 14th International Conference (DEON)*, ed. Jan Broersen, Cleo Condoravdi, Shyam Nair, and Gabriella Pigozzi, 227–246. London: College Publications.
- Mandelkern, Matthew, Ginger Schultheis, and David Boylan. 2017. “Agentive Modals.” *Philosophical Review* 126(3): 301–343. <http://dx.doi.org/10.1215/00318108-3878483>.
- Marušić, Berislav. 2015. *Evidence and Agency: Norms of Belief for Promising and Resolving*. New York: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780198714040.001.0001>.
- Meinong, Alexius. 1917. *Über Emotionale Präsentation*. Wien: A. Hölder.
- van der Meyden, Ron. 1996. “The Dynamic Logic of Permission.” *Journal of Logic and Computation* 6(3): 465–479. <http://dx.doi.org/10.1093/logcom/6.3.465>.
- Meyer, John-Jules Ch. 1987a. “A Different Approach to Deontic Logic: Deontic Logic Viewed as a Variant of Dynamic Logic.” *Notre Dame Journal of Formal Logic* 29(1): 109–136. <http://dx.doi.org/10.1305/ndjfl/1093637776>.
- . 1987b. “A Simple Solution to the “Deepest” Paradox in Deontic Logic.” *Logique et Analyse* 30(117/118): 81–90. <https://www.jstor.org/stable/44084186>.
- Meyer, John-Jules Ch., Wiebe van der Hoek, and Bernd van Linder. 1999. “A Logical Approach to the Dynamics of Commitments.” *Artificial Intelligence* 113(1-2): 1–40. [http://dx.doi.org/10.1016/S0004-3702\(99\)00061-2](http://dx.doi.org/10.1016/S0004-3702(99)00061-2).
- Ninan, Dilip. 2005. “Two Puzzles about Deontic Necessity.” In *MIT Working Papers on Linguistics 51: New Work On Modality*, ed. Jon Gajewski, Valentine Hacquard, Bernard Nickel, and Seth Yalcin, 149–178. Cambridge, MA: MIT Press.

- Nouwen, Rick. Forthcoming. “Free Choice and Distribution over Disjunction: The Case of Free Choice Ability.” *Semantics and Pragmatics*. <http://dx.doi.org/10.3765/sp.11.4>.
- Pietroski, Paul M. 1998. “Actions, Adjuncts, and Agency.” *Mind* 107(425): 73111. <http://dx.doi.org/10.1093/mind/107.425.73>.
- Portner, Paul. 2009. *Modality*. New York: Oxford University Press.
- Prior, Arthur N. 1967. *Past, Present and Future*. Oxford: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780198243113.001.0001>.
- Rooth, Mats, and Barbara H. Partee. 1983. “Generalized Conjunction and Type Ambiguity.” In *Meaning, Use and Interpretation of Language*, ed. Rainer Bäuerle, Christoph Schwarze, and Armin von Stechow, 361–383. Berlin: De Gruyter.
- Ruwet, Nicolas. 1984. “Je veux partir / *Je veux que je parte: On the Distribution of Finite Complements and Infinitival Complements in French.” *Cahiers de Grammaire* 7: 75–138.
- Schlenker, Philippe. 2005. “The Lazy Frenchman’s Approach to the Subjunctive.” In *Romance Languages and Linguistic Theory 2003: Selected papers from ‘Going Romance’ 2003, Nijmegen, 20-22 November*, ed. Twan Geerts, Ivo van Ginneken, and Haïke Jacobs, 269–309. Amsterdam: John Benjamins. <http://dx.doi.org/10.1075/cilt.270.15sch>.
- Schroeder, Mark A. 2011. “Ought, Agents, and Actions.” *Philosophical Review* 120(1): 1–41. <http://dx.doi.org/10.1215/00318108-2010-017>.
- Schulz, Katrin. 2005. “A Pragmatic Solution for the Paradox of Free Choice Permission.” *Synthese* 147(2): 343–377. <http://dx.doi.org/10.1007/s11229-005-1353-y>.
- Seegerberg, Krister. 1982. “A Deontic Logic of Action.” *Studia Logica* 41(2-3): 269–282. <http://dx.doi.org/10.1007/BF00370348>.
- . 1989. “Bringing It about.” *Journal of Philosophical Logic* 18(4): 327–347. <http://dx.doi.org/10.1007/BF00262940>.
- . 1992. “Action Incompleteness.” *Studia Logica* 51(3/4): 533–550. <http://dx.doi.org/10.1007/BF01028974>.
- Sergot, Marek. 2012. “Normative Positions.” In *Handbook of Deontic Logic and Normative Systems*, ed. Dov Gabbay, John Horty, Xavier Parent, Ron van der Meyden, and Leon van der Torre, 353–406. London: College Publications.
- . 2014. “Some Examples Formulated in a ‘Seeing to It That’ Logic: Illustrations, Observations, Problems.” In *Nuel Belnap on Indeterminism and Free Action. Outstanding Contributions to Logic, vol 2.*, ed. Thomas Müller, 223–256. Cham: Springer.
- Silk, Alex. 2015. “How To Be an Ethical Expressivist.” *Philosophy and Phenomenological Research* 91: 47–81. <http://dx.doi.org/10.1111/phpr.12138>.

- . Forthcoming. “Weak and Strong Necessity Modals.” In *Meaning, Decision, and Norms: Themes from the Work of Allan Gibbard*, ed. Bill Dunaway and David Plunkett. Ann Arbor, MI: Michigan Publishing.
- Simons, Mandy. 2005. “Dividing Things up: The Semantics of *Or* and the Modal/*Or* Interaction.” *Natural Language Semantics* 13(3): 271–316. <http://dx.doi.org/10.1007/s11050-004-2900-7>.
- Starr, William B. 2016. “Expressing Permission.” In *Proceedings of SALT XXVI*, ed. Mary Moroney, Carol-Rose Little, Jacob Collard, and Dan Burgdorf, 325–349. Ithaca, NY: CLC Publications. <http://dx.doi.org/10.3765/salt.v26i0.3832>.
- Thomason, Richmond H. 1970. “Indeterminist Time and Truth-Value Gaps.” *Theoria* 36(3): 264–281. <http://dx.doi.org/10.1111/j.1755-2567.1970.tb00427.x>.
- . 1981. “Deontic Logic and the Role of Freedom in Moral Deliberation.” In *New Studies in Deontic Logic*, ed. Risto Hilpinen, 177–186. Dordrecht: Reidel. http://dx.doi.org/10.1007/978-94-009-8484-4_8.
- . 2005. “Ability, Action, and Context.” <http://web.eecs.umich.edu/~rthomaso/documents/action/ability.pdf>. Manuscript, University of Michigan.
- Veltman, Frank. 1985. *Logics for Conditionals*. Ph.D. thesis, University of Amsterdam.
- . 1996. “Defaults in Update Semantics.” *Journal of Philosophical Logic* 25(3): 221–261. <http://dx.doi.org/10.1007/BF00248150>.
- Wedgwood, Ralph. 2006. “The Meaning of ‘Ought’.” In *Oxford Studies in Metaethics*, ed. Russ Shafer-Landau, vol. 1, 127–160. New York: Oxford University Press.
- . 2007. *The Nature of Normativity*. New York: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199251315.001.0001>.
- Willer, Malte. 2015. “Simplifying Counterfactuals.” In *Proceedings of the 20th Amsterdam Colloquium*, ed. Thomas Brochhagen, Floris Roelofsen, and Nadine Theiler, 428–437. Amsterdam: ILLC Publications.
- . 2017. “Widening Free Choice.” In *Proceedings of the 21st Amsterdam Colloquium*, ed. Alexandre Cremers, Thom van Gessel, and Floris Roelofsen, 511–520. Amsterdam: ILLC Publications.
- . 2018. “Simplifying with Free Choice.” *Topoi* 37(3): 379–392. <http://dx.doi.org/10.1007/s11245-016-9437-5>.
- Williams, Bernard. 1991. “Moral Luck.” In *Ought and Obligation*, ed. Bernard Williams, 114–123. Cambridge: Cambridge University Press. <http://dx.doi.org/10.1017/CB09781139165860.010>.
- von Wright, George H. 1951. “Deontic Logic.” *Mind* 60(237): 1–15. <http://dx.doi.org/10.1093/mind/LX.237.1>.
- . 1968. *An Essay on Deontic Logic and the Theory of Action*. Amsterdam: North-Holland.

Zimmermann, Thomas Ede. 2000. "Free Choice Disjunction and Epistemic Possibility." *Natural Language Semantics* 8(4): 255–290. <http://dx.doi.org/10.1023/A:1011255819284>.