Ability as Dependence Modality [1]

Paolo Santorio

University of Maryland, College Park

Penultimate version. Final version forthcoming in *Noûs*.

1 Introduction

Modal expressions in language can describe what is possible in light of a subject's abilities. In English, modals of this sort include the modal auxiliary *can*, as well as the predicate *able*. Here are some examples:

(1) a. Ava can hit the target on the next throw.
b. Ben is able to join the conference virtually.
c. Clem can run 100m in 10 seconds.

Ability modals are obviously related to other modalities in language, such as epistemic or deontic modality, but also give rise to anomalies that make them unique.

This paper develops a general theory of ability modals that is broadly compatible with standard modal semantics, while predicting their peculiar behavior. The central idea is that ability modals include reference to a notion of dependence. Roughly, (1a) requires that there is an accessible world where Ava hits the target, and that Ava's hitting the target depends on features of Ava, in some relevant sense of dependence.

The appeal to dependence is idiosyncratic to these modals, but the implementation is compatible with a classical semantic architecture. At the basic level, ability modals are simply existential quantifiers over worlds. In addition, they include a not-at-issue element in their meaning that enforces the

[1]Thanks to Harjit Bhogal, David Boylan, Fabrizio Cariani, Ilaria Canavotto, Cleo Condoravdi, Melissa Fusco, Matt Mandelkern, Dilip Ninan, Alexander Williams, and audiences at the University of Maryland, the Dublin Language Workshop, and the Amsterdam Colloquium. Special thanks to Ilaria Canavotto and Fabrizio Cariani for extensive feedback at all stages of the project.

dependence requirement. The latter is what gives rise to the apparent logical anomalies exhibited by *can* and *able*.

The dependence analysis is indebted to some previous and modern accounts in the literature (in particular, Hackl 1998, and less directly Kenny 1976; see also Nadathur 2021[2]). At the same time, it contrasts with most existing accounts. In particular, it divorces ability modals from a notion of agency. Several accounts appeal to agential notions to analyze the meaning of *can* and *able*. This connection is developed in various ways, but some ideas are recurring: (i) the subject of ability modals is an agent; (ii) ability reports involve quantification over actions; (iii) ability reports require that the subject have, in some sense, control over the event denoted by the complement clause. All of (i)–(iii) are incorrect, at least if we're concerned with the concept of ability that is encoded in language.

I proceed as follows. I provide some background in §2. In §3, I introduce the main empirical puzzle about ability modals: they are stronger than their circumstantial counterparts in both positive and negative contexts. I state my positive theory in §5 and §6; finally, in §7 I show that the agentive analysis is unable to explain some crucial facts.

2 Ability modals: basic facts

In this section, I outline some general assumptions that will work as background for my discussion. These assumptions are not uncontroversial, but they are widely adopted.

It is well-know that the same modal items can express different modal flavors (Kratzer 1977, 1981, 2012). For example, as illustrated by (2), *must* can have epistemic, deontic, or circumstantial flavor—expressing what is necessary in light of a body of evidence, a set of norms, or a set of facts, respectively.

(2) a. Acacia must be out of her office.
b. Bashir must submit the forms by tomorrow at noon.
c. Cody must sneeze.

[2]Interestingly, Nadathur uses causal notions to solve puzzles related to ability modals, and in particular actuality entailments; but she still adopts an agential analysis for the lexical semantics of ability modals.

A first assumption driving this paper is that ability is an additional **dedicated modal flavor**. *Able* invariably expresses this flavor; *can* expresses this flavor in some of its uses. I will assume that *S is able to A* and *S can A*, when *can* has an ability use, are synonymous.

In §3 I will discuss ability modals *vis-á-vis* circumstantial modals. So let me also introduce the latter. Circumstantial modals express what is possible or necessary in light of certain relevant facts.[3] A classical example of a circumstantial modals is in (3) (adapted from Kratzer 1981).

(3) Hydrangeas can grow in this region.

(3) can be true even in a scenario where a speaker knows that no hydrangeas grow in the region. In this scenario, it means that, in light of some relevant facts (presumably having to do with water, light, soil, etc.) it is possible that hydrangeas grow there. As I will point out, ability modals are related to circumstantials, but diverge from them in significant ways.

The second assumption is that ability modals have **two kinds of occurrences, specific and generic**. Specific occurrences, exemplified by (4a), describe a specific event that is possible in the light of the subject's abilities. Generic occurrences, exemplified by (4b), describe what is possible in general for a subject, given their abilities.

(4) a. Ava is able to hit the target on this throw.
b. Ava is able to hit the target when playing darts.

Given that natural language involves a generic operator GEN (see e.g. Leslie & Lerner 2022), the natural way to accommodate this fact is to assume that the basic reading of ability modals is the specific one. Generic readings are derived from the specific one by combining ability modals with GEN. This is the strategy that I adopt in this paper, hence my focus will be on specific readings. (See Mandelkern et al. 2017 for an analogous strategy.) Notice that also the discussion of the data, in §3, only concerns specific readings.

The third assumption has to do with the argument type of ability modals. On standard semantic analyses derived from modal logic, modals are sentential operators. This means that their syntactic argument is a full clause, which semantically denotes a proposition. For the case of ability modals, it's unclear that this analysis is correct. A different analysis might be plausible:

[3] This characterization is vague, but so is, notoriously, the category of circumstantial modals. For discussion, see Kratzer 1981.

can and *able* take two arguments, a subject and a verb phrase-type argument, which denote respectively an individual and an event, or a property.[4]

Be that as it may, this issue is not central for my purposes, so I set it aside. I make the simplifying assumption that ability modals are syntactically on a par with other modals, and take a sentential complement. Following current practice, I call this complement *prejacent*.

There is a further complication, related to tense. Strictly speaking, prejacents of ability modals are tenseless phrases. E.g., the prejacent of *Ava is able to hit the board* is *Ava hit the board*. So the clauses that appear embedded under ability modals never appear as independent utterances, but they are always combined with tense, or with *will*.[5] So far as I can see, this makes no difference to my arguments. So I make the simplifying assumption that the prejacents of ability modals are (equivalent to) utterances involving tense or *will*. So e.g. the prejacent of (5b) is (equivalent to) (5a).

(5) a. Ava will hit the board.
    b. Ava can hit the board.

3 The exceptional strength of ability

On standard theories of modality (Kratzer 1977, 1981, 2012), modal expressions of different flavors share a logical skeleton, which is analyzed in terms of quantification over possible worlds. The driving idea is that each modal

[4]One fact that might support this view is that ability modals are, differently from e.g. epistemic modals, control predicates. (The distinction between control and raising predicates is standard in syntax; , see e.g. Carnie 2021.) The notion of a control predicate is a syntactic notion, but it is plausible that it has implications for the semantics. Among other things, the fact that ability modals are control predicates means that they don't have semantic interactions with quantifiers. (1) has two readings, depending on the relative scope of the determiner phrase *most students* and the modal *allowed*. (On the first reading, there is a majority of students such that each of them has individual permission to enroll; on the second reading, it is allowed that a majority of students enroll in the class.) Not so for (2).

(1) Most students are allowed to enroll in this class this semester.

(2) Most students are able to enroll in this class this semester.

This suggests that *able* takes a separate subject argument.

[5]For arguments that *will* is not a tense see, among many, Klecha 2013, Cariani & Santorio 2018, Cariani 2021.

expression only has one core meaning. Differences in modal flavor are captured via contextual parameters.

The natural starting assumption for a semantics of ability modals is that they conform to this model. In particular, we should expect them to be possibility modals, i.e. existential quantifiers over a domain of worlds. Hence the meaning for *able* would be stated as follows:

(6) ⌜*S* is able to A⌝ is true at *w* just in case there is a world *w′*, accessible from *w*, such that ⌜*S* As⌝ is true in *w′*.

This section show that a semantics along the lines of (6) is untenable. I investigate systematically the relationship between ability and circumstantial modals. Some of the evidence that I review is known, but some is not, and the picture that I arrive at is altogether new.

To preview: I argue that *able* entails circumstantial *can*, and *not able* entails circumstantial *cannot*. This kind of behavior requires a nonclassical treatment. I will suggest that the best way to account for it is to postulate that ability modals have a not-at-issue element in their meaning—possibly, though not necessarily, a presupposition.

Before I start, a word of clarification. I present my evidence as concerning patterns of entailment between sentences. This is intended in a purely descriptive way. By saying that A entails B, I mean simply that the inference from A to B is licensed in virtue of facts about the meaning of A and B. This is compatible with the claim that, ultimately, the logic of these sentences should be systematized via a notion of consequence that does not validate an entailment from A to B.[6]

3.1 Ability and circumstantial possibility

It is well known that ability modals are stronger than their circumstantial counterparts. Consider a scenario where Ava is a first time dart thrower and has no other related skills. An assertion of (7a) would be felicitous, but an assertion of (7b) would not.

[6]For example, some theorists rely on a multidimensional view of presupposition, on which sentences that involve presupposition failure can still be true or false (see Sudo 2012, Mandelkern 2023; see Herzberger 1973 for a classical explorations of these ideas). These views allow one to define a classical notion of entailment even in the presence of presupposition. Nothing that I say is in principle incompatible with adopting this notion of consequence.

(7) a. Ava can hit the target on this throw—sometimes beginners are lucky!

b. # Ava is able to hit the target on this throw—sometimes beginners are lucky!

(7a) has a true reading. This reading is synonymous with the sentence we get by using an expletive subject, as in (8):

(8) It can happen that Ava hits the target on this throw.

(8) exemplifies the circumstantial reading of the modal. On this reading, *can* does not have to do with Ava's abilities. Rather, *can* tracks what is possible, in light of certain relevant facts. Conversely, (7b) does not have a true reading in this scenario. (Notice: I don't take this to mean that (7b) is *false*. More below.) This is evidence that circumstantial possibility does not entail the modality expressed by *able*.

Another example showing that ability modals are stronger than circumstantials is due to Wolfgang Schwarz (2020). Schwarz imagines a scenario where a statue is precariously placed on a ledge. He notices that (9a) is true, but (9b) is not, and moreover sounds very awkward.

(9) a. The statue can [easily] fall. (Schwarz 2020, p. 3)

b. # The statue is able to [easily] fall.

In sum: there is a failure of entailment from the claim that it's circumstantially possible that *S* As, to the claim that *S* is able to A. What about the converse? A quick survey of examples suggests that this direction of the entailment holds.

(10) a. Ava is able to hit the target on this throw.
$\rightsquigarrow$ It can happen that Ava hits the target on this throw.

b. Ben is able to join the conference virtually.
$\rightsquigarrow$ It can happen that Ben joins the conference virtually.

c. Clem is able to run 100m in 12 seconds in this trial.
$\rightsquigarrow$ It can happen that Clem runs 100m in 12 seconds in this trial.

Indeed, it seems obvious that ability should entail circumstantial possibility. If *S* is able to A, there has to be a possibility where *S* As.[7] Unsurprisingly, this entailment is validated by pretty much all existing semantics for ability.

[7] Though see Spencer 2017 for a dissenting view.

To sum up: using the dotted diamond '$\diamondsuit\!\!\!\cdot$' to model ability modals and the simple diamond '$\Diamond$' to model circumstantial modals, we have:

Ability asymmetrically entails circumstantial possibility.

$\diamondsuit\!\!\!\cdot A \vDash \Diamond A$	*S* is able to A $\vDash$ *S* can A
$\Diamond A \nvDash \diamondsuit\!\!\!\cdot A$	*S* can A $\nvDash$ *S* is able to A

3.2 Circumstantial impossibility and inability

The foregoing may suggest that ability is simply a stronger notion than circumstantial possibility—in the same way as, say, physical possibility is a stronger notion than metaphysical possibility. But the picture becomes more complex once we consider negation.

Take again the case of Ava, the first time dart thrower, and consider:

(11) # Ava is not able to hit the target on this throw.

(11) is the negation of (7b). And like (7b), it seems defective in the context described. Ava might get lucky and hit the target on her first throw. This is enough to make (11) unassertable, on a par with (7b).

A second relevant observation is that the truth of *S is not able to A* is incompatible with the circumstantial possibility of A, as shown by (12):

(12) # Ava is not able to hit the dartboard on this throw, but it can happen that she hits the dartboard on this throw (by sheer luck).

Assuming that circumstantial modals behave classically with respect to negation, this means that *S is not able to A* entails *It cannot be that S As*. This entailment is also confirmed by a quick survey of simple cases.

(13) a. Xavier is not able to play the piano for us tomorrow.
$\rightsquigarrow$ It cannot happen that Xavier plays the piano for us tomorrow.

b. Yuri is not able to win this match.
$\rightsquigarrow$ It cannot happen that Yuri wins this match.

c. Zina is not able to run 100m in 10 seconds in this trial.
$\rightsquigarrow$ It cannot happen that Zina runs 100m in 10 seconds in this trial.

Finally, it's easy to see that the converse entailment does not hold. *S cannot A* may be true, and yet *S is not able to A* may fail to be true. To see this, just consider a variant of Schwarz's example. The statue has now been anchored

very solidly to its ledge. (14a) is true, but (14b) is not, and again sounds very awkward.

(14) a. The statue cannot fall from the ledge.
b. #The statue is not able to fall from the ledge.

To sum up, the following two principles also seem to hold:

Inability asymmetrically entails circumstantial impossibility.

$\neg \diamondsuit \mathsf{A} \vDash \neg \Diamond \mathsf{A}$	*S* is not able to A $\vDash$ *S* cannot A
$\neg \Diamond \mathsf{A} \nvDash \neg \diamondsuit \mathsf{A}$	*S* cannot A $\nvDash$ *S* is not able to A

3.3 The mark of nonclassicality

The fact that ability modals validate this array of principles is puzzling. Take the principle that ability entails circumstantial possibility. By contraposition, we get that circumstantial impossibility entails inability.

$$\diamondsuit \mathsf{A} \vDash \Diamond \mathsf{A} \quad \Rightarrow \quad \neg \Diamond \mathsf{A} \vDash \neg \diamondsuit \mathsf{A}$$

Yet we just established that this principle is invalid. Similarly for the principle that inability entails circumstantial impossibility. By contraposing, we get that circumstantial possibility entails ability. But classical notions of consequence validate contraposition. So, if the findings of this section are correct, the relation of entailment we're tracking must be nonclassical.

In this paper, I suggest that this nonclassical behavior should be captured by postulating a not-at-issue element in ability modals. Not-at-issue meaning is a well-known phenomenon in semantics. Some parts of the meaning of a sentence appear to be 'backgrounded', and work differently in compositional processes. The classical example of not-at-issue meaning in natural language is presupposition. Items that have presuppositions include definite descriptions, factive verbs like *know*, and cleft-type constructions like *S is the first to A*.

(15) a. Ava's personal coach is Iranian. $\rightsquigarrow$ Ava has a personal coach.
b. Ben knows that Clem won. $\rightsquigarrow$ Clem won.
c. D'Auria is the first to get a perfect score. $\rightsquigarrow$ D'Auria got a perfect score.

Presuppositions project along systematic patterns in complex sentences.[8] Among other things, they survive under negation: if S presupposes p, $\neg S$ also presupposes p. This generates nonclassical effects in the logic. In particular, a notion of entailment that tracks presupposition may fail to contrapose.[9] For an illustration, notice that the entailment in (16a) holds, but the one in (16b) obviously doesn't.

(16) a. Ben doesn't knows that Clem won $\models$ Clem won
b. Clem didn't win $\not\models$ Ben knows that Clem won

Presupposition is the best known type of not-at-issue meaning that shows these projection properties under negation, but not the only one. In this paper, I remain neutral about exactly what kind of not-at-issue meaning is involved in ability modals. What matters is that this component is there, and induces nonclassical features in the logic of ability modals.

4 Overview: ability and dependence

In §3, I suggested that ability modals include an extra, not-at-issue component in their meaning. In this section, I explain intuitively what this extra element says.

Consider again an ability ascription, and compare it to the corresponding circumstantial possibility claim:

(17) a. Ava is able to hit the target on this throw.
b. It can happen that Ava hits the target on this throw.

I suggest that (17a), but not (17b), requires that whether Ava hits the target *depends on Ava*, as opposed to luck or external circumstances of various kinds. This dependence claim is understood as a sufficiency claim: some relevant facts about of Ava (plus, as we'll see, some background facts) determine whether or not she hits the target.[10]

[8] For some classical references, see Karttunen 1974, Heim 1983, Schlenker 2008.

[9] One standard way to formulate a notion of entailment that tracks presupposition is Strawson-Entailment; see von Fintel 1999 for a classical reference. See Cariani & Goldstein 2018 for the explicit point that Strawson-Entailment has some nonclassical features. Strawson-Entailment in particular does contrapose, but it's easy to define a variant notion that invalidates contraposition.

[10] Hence some commonly used notions of dependence in the literature, such as counterfactual dependence, are much weaker than the notion I'm interested in. This seems correct: even in a case where Ava hits the target by chance, her hitting the target counterfactually depends on her throw.

To motivate the view intuitively, let me introduce two examples.[11]

> **Magical dart.** Ben is a mediocre dart thrower who's about to throw a dart. In ordinary circumstances, there would be a high chance that he would miss. But Ben's magician friend Camille wants Ben's dart to hit the target. So, as soon as the dart leaves Ben's hands, Camille will cast a spell on the dart, leading it to the target.

Notice first that, in this scenario, the circumstantial necessity claims in (18) are true. (I take the claims to be truth-conditionally equivalent.)

(18) a. It cannot happen that Ben doesn't hit the target on this throw.
b. Ben cannot miss the target on this throw.

Yet there is at least one salient reading on which (19) doesn't sound true.

(19) # Ben is able to hit the target on this throw.

Why isn't (19) true? Intuitively, the problem is that anything that Ben does is irrelevant to whether he hits the target. Ben's hitting the target does not depend on Ben, at least not in any relevant way.

And now consider:

> **The baby carrier.** Someone is carrying their infant daughter in a baby carrier. The baby is leaning out in a way that appears dangerous, and you worry that she might fall. But the carrier is actually very safe.

The discourse in (18) perfectly felicitous, and in particular the circumstantial necessity claim in (19) has an obvious true reading.

(20) a. Don't worry! The carrier is very secure...
b. The baby cannot fall.

Conversely, both (21b) and (21c) are very awkward.

(21) a. Don't worry! The carrier is very secure...
b. # The baby is not able to fall.
c. # The baby is able to not fall.

[11] The first of these examples is a variant of Lewis's (1997) well-known example of the sorcerer and the glass (see §5). The second example is obviously an elaboration of Schwarz's statue example.

Intuitively, the problem with these sentences is that whether the baby falls or not does not depend on the baby herself. No feature or action of the baby will make a difference to her falling. Whether she falls entirely depends on external factors—i.e., the carrier.

In the next two sections, I develop a view of ability modals that revolves around this idea.[12] I spell out the general account in §5, and implement it in a formal semantics in §6. Examples that motivate the dependence idea will also serve to refute the idea that ability is linked to agency, in §7.

5 The dependence analysis

I suggest that, in addition to the usual quantification over possibilities, ability claims encode a dependence requirement. *S is able to A* requires that whether *S* As fully depends on features of *S*, plus laws of nature and some background facts. In other words: features of *S*, together with background facts, are sufficient to determine that A, or sufficient to determine that not A.

The account, then, exploits notions of sufficiency and dependence. One natural thought is that these notions should be understood causally. This thought is very much on the right track, but some ability reports resist a causal analysis. For some examples, consider:[13]

(22) The theory is able to explain all the data.

(23) This algorithm is able to approximate the solutions of these differential equations.

(24) This logic is able to validate all theorems of propositional logic.

So, while the relevant notion of dependence can be understood as causal in a great variety of cases, I state the view in more general terms. This said, all

[12] The view I develop is obviously related to a view on which ability reports describe powers, or potentialities, of their subjects. For a view on which modal claims in general describe powers or potentialities, see Barbara Vetter's work on modality (2015). Notice that, in the metaphysics literature, powers and potentialities are linked to non-Humean accounts of causation and lawhood, on which just the notion of power is an unreduced primitive (besides Vetter, see e.g. Bird 2007). My account is neutral on issues of metaphysical reducibility. All I need is that intrinsic features of the subject of ability ground some relevant dependencies.

[13] Thanks to an anonymous referee, as well as to *[name omitted for anonymous review]* for pressing me on this.

my core examples will be causal.[14]

The plan for this section is the following. First, I characterize what dependencies are relevant for ability (§5.1). Then I show how we can characterize notions like dependence and sufficiency in a rigorous way (§5.2). Then I state informal truth conditions for ability ascriptions (§5.3) and briefly discuss the role of context (§5.4).

5.1 Ability and dependencies

So dependencies are relevant for evaluating ability claims. Which dependencies? I suggest: dependencies that are grounded in intrinsic properties of the subject of ability. For example, to evaluate ability claims about Ava, we consider possibilities where Ava has all her actual intrinsic properties, but external circumstances may be different. This builds on ideas that have been developed by Lewis in his analysis of dispositions (1997), and exported by Kratzer (2013) to semantics.

Start by considering a well-known example from Lewis's (1997) discussion of dispositions:

> A sorcerer takes a liking to a fragile glass, one that is a perfect intrinsic duplicate of all the other fragile glasses off the same production line. He does nothing at all to change the dispositional character of his glass. He only watches and waits, resolved that if ever his glass is struck, then, quick as a flash, he will cast a spell that changes the glass, renders it no longer fragile, and thereby aborts the process of breaking. So his finkishly fragile glass would not break if struck—but no thanks to any protective disposition of the glass itself. Thanks, instead, to a disposition of the sorcerer.

Lewis notices that the following is clearly true:

(25) The glass is fragile.

Nevertheless, given the situation, the following circumstantial impossibility claim is also true:

(26) The glass cannot break.

[14] In fact, capturing all cases of dependence, including cases of logical and mathematical dependence as in (23) and (24), will require dropping the standard possible worlds framework in favor of a framework that is friendlier to hyperintensional notions. This generalization goes way beyond the remit of this paper, of course.

(25) is, on a natural analysis, a modal claim. It says that, in some relevant sense of possibility, it's possible (in fact, easily possible) that the glass breaks. So its meaning involves reference to a set of possibilities. Which set? Evidently, a set of possibilities that is not a subset of the set of circumstantially accessible worlds, since the glass stays intact throughout the latter (as witnessed by the truth of (26)). To generate this domain of worlds, it appears that we do the following: we hold fixed all the intrinsic features of the glass—basically, all facts about its physical composition and molecular structure—but we allow changes in external circumstances. In particular, we consider worlds where there is no sorcerer, and hence the glass breaks upon being struck. (Notice that we also hold fixed some other facts external to the glass, such as physical laws.)

More recently, Kratzer (2013) has used examples like (25) to defend a general view about the metasemantics of modality. She suggests that modal domains are 'projected' from objects and events that are mentioned in the clause. She calls the relevant objects 'anchors' and suggests that modal domains in general are generated via anchoring.[15] Lewis's glass example is just an example of modal anchoring.

Here I remain neutral on the claim that modal domains in general are determined via anchoring. But I suggest that ability modals are an example of modal anchoring in Kratzer's sense: they exploit a domain of possibilities that is generated by holding fixed all intrinsic features of the subject of ability. Consider again:

(27) Ava is able to hit the target on this throw.

I suggest that, to evaluate (32), we consider a domain of worlds D—the **dependence domain**—with the following features.

(i) Worlds in D agree with the actual world with respect to Ava's intrinsic properties: Ava figures throughout D as she actually is.

(ii) Worlds in D validate physical laws.

[15]Kratzer builds on work on modality produced in the previous ten years, in particular Arregui 2004, 2007, 2009 and Hacquard 2006, 2010. I should emphasize that Kratzer might not agree with my use of modal anchors here. In particular, she seems to suggest that the domains fixed by modal anchors should altogether replace domains fixed via the so-called 'conversational backgrounds' in Kratzer's own classical theory (i.e. modal base and ordering source). My semantics still appeals to a modal base and an ordering source.

(iii) Worlds in *D* validate some other array of propositions, fixed by context. These include the proposition that the subject's abilities are triggered, in relevant ways. For example, in Ava's case, they include the proposition that she has an intention to throw the dart.

Importantly, worlds in the dependence domain can overlap with circumstantially accessible worlds, but they can also form a disjoint set. That they can be disjoint is shown by cases that are analogous to Lewis's sorcerer scenario: these are cases where *S* A-ing is circumstantially necessary, yet the claim that *S* is able to A is not true. I already presented a case of this sort, i.e. Ben's case from §4. Recall: Ben is a mediocre dart thrower, but his magician friend Camille will guide his dart to the target. In this scenario, (18b) is true, but (19) has at least one reading on which it's not.

(18b) Ben cannot miss the target on this throw.

(19) Ben is able to hit the target on this throw.

This means that, even if Ben hits the target in all circumstantially accessible worlds, when evaluating (19) we still consider possibilities where Ben doesn't hit. In these possibilities we keep fixed facts about Ben, but not facts about Camille. These are the possibilities in the dependence domain. (This said, while the two sets can be disjoint, they will always overlap when an ability ascription is true. I say more in §5.3.)

The dependence domain is used to check that the relevant dependencies hold, i.e., that the features of the subject of ability, together with the assumed background, determine whether the prejacent of *able* holds. For illustration, consider again (32).

(32) Ava is able to hit the target on this throw.

Suppose that Ava, who has the brain and muscles of an expert dart thrower, is about to throw a dart. The dependence domain for (32) will include worlds where (i) Ava has the same intrinsic properties she actually has, and hence where her throwing skills hold; (ii) the laws of physics obtain; (iii) some relevant background facts obtain (the target is at the same distance as in the actual world, there won't be sudden gusts of wind, etc), and moreover Ava has the intention to throw. The dependence requirement is satisfied just in case, in all these worlds, Ava hits the target when throwing.

5.2 Dependence and sufficiency: a formal characterization

So far, I have relied on intuitive notions of dependence and sufficiency. This section provides a more precise characterization of these notions, building on tools available in the literature. Readers who are happy with an informal characterization can skip ahead.

Dependence and sufficiency can be characterized precisely in many ways. For concreteness, here I provide a characterization in terms of causal models (Pearl 2000, Halpern 2000, Halpern 2016). I appeal to causal models because they are intuitive and popular, and because their connection with semantics has been studied already (see e.g. Kaufmann 2013, Santorio 2019). Moreover, as the philosophical literature has pointed out, causal models can be easily generalized to noncausal cases (see e.g. Schaffer 2016, Wilson 2018). But the basic idea behind my account, as well as the general semantic framework in §6, can be combined with other accounts of dependence.

A causal model consists of three elements: a set of **random variables**, a set of **structural equations**, and an **assignment of values** to the variables. A **random variable** can be thought of as a set of mutually exclusive and jointly exhaustive outcomes for a process. Following standard notation, I represent random variables with uppercase Roman letters, like 'X'. I represent the values of random variables with lowercase Roman letters, sometimes with subscripts, like 'x_i'. Moreover, to say that variable X has value x_i, I use the notation '$X = x_i$'. An equation like '$X = x_i$' basically says that a certain event obtains. (For example, using variables from the model below, we can represent the event of a target being hit with the equation '$T = 1$'.)

Structural equations are mathematical equations that state the relations between different values of random variables. An **assignment of values** is simply a mapping of each variable in a model to one of its possible values.

Let me go through an example in detail. Consider once more Ava the dart thrower, in front of the target and about to throw. The factors that determine whether Ava hits are: (i) Ava's throwing skills (which, on the present construal, are fixed by Ava's intrinsic properties); (ii) a series of background circumstances, such as whether there are gusts of wind, whether someone pushes Ava while throwing, etc; (iii) possibly, chance. (Chance will be relevant in case Ava's throwing skills are insufficient to determine the outcome of the throw.)

We can represent this situation with a model that includes six variables:

B : whether background circumstances (wind, etc) are favorable
C : whether chance goes Ava's way
I : whether Ava has an intention to throw
S : whether Ava is a skilled thrower
T : whether Ava throws well
H : whether Ava hits the target

I assume that all variables are binary (i.e. they have values 0 or 1); more sophisticated models will have variables with a wider range of values.

Let us move to the equations. Whether Ava hits the target depends on whether she throws well, plus background circumstances. Whether Ava throws well depends on her intention to throw, on her throwing skill, and possibly on chance. This is captured by the following equations:

T = min(I, max(S, C))
H = min(B, T)

The first equation says that Ava throws well iff it is both the case that she intends to throw and that either she is skilled, or she gets lucky. The second equation says that Eva hits the target if she throws well and the background circumstances are right. These equations can be represented by the following causal diagram:

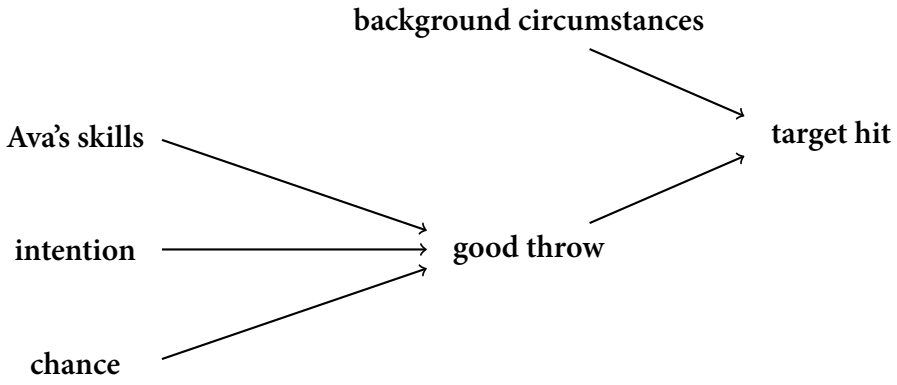

Figure 1. Toy causal model for the dart throw scenario.

I'm construing causal models as involving also a specification of values for the variables. To get a complete model, let's assume that Ava intends to throw ($I = 1$), that she is a skilled dart thrower ($S = 1$), that she is not particularly lucky ($C = 0$), and that the background circumstances are favorable

($B = 1$). It's easy to check that variables T and H get value 1 and that hence Ava hits the target.

Now, against this background, I define sufficiency as follows:

Sufficiency. An event $X = x$ is **sufficient** for an event $E = e$, relative to a set of equations $\mathcal{S}$ and background circumstances $\mathcal{C} = \{C_1 = c_1, \ldots, C_n = c_n\}$ iff any model that is (i) consistent with the equations in $\mathcal{S}$, (ii) consistent with the background circumstances $\mathcal{C}$, and (iii) such that $X = x$, is also a model where $E = e$.

Informally, the idea is simple. Sufficiency is a kind of necessitation. An event is sufficient for an effect e just in case, holding fixed the equations and some background circumstances, any model compatible with the equations where the event obtains is also a model where e obtains.[16]

We can define dependence on the basis of sufficiency. I construe dependence as a relation between a set of outcomes (which for simplicity I take to be mutually exclusive) and an event.

Dependence. A set of jointly exhaustive and mutually exclusive events $\{E = e_1, \ldots, E = e_n\}$ **(fully) depends** on an event $X = x$, relative to a set of equations $\mathcal{S}$ and background circumstances $\mathcal{C} = \{C_1 = c_1, \ldots, C_n = c_n\}$ iff, for some $E = e_i$ in the set, $X = x$ is sufficient for $E = e_i$.

Intuitively: a set of alternative outcomes depends on an event just in case the event is sufficient for one of the outcomes. (As the parentthetical makes clear, by 'dependence' here I mean a notion of full dependence.)

Let me illustrate how this characterization handles some paradigm examples. Consider first the Ava case, already diagrammed in Figure 1. In that case, we have that whether Ava hits depends on the fact that Ava is a skilled thrower. Holding fixed the model (i.e. the structural equations), the background facts, and the fact that Ava intends to throw, the fact that Ava is skilled ($S = 1$) is sufficient for Ava hitting ($T = 1$).

Conversely, consider Ben, a less skilled dart thrower. We can model Ben's situation with a causal model that is fully analogous to the model for Ava, aside from the fact that the S variable gets a different value ($S = 0$). In this case, Ben's skills, holding fixed all the rest, are not sufficient to determine

[16]This definition is reminiscent of the so-called 'INUS' account of causation, first introduced by Mackie 1965 and then elaborated by many authors.

the fact that he hits the target. Chance (i.e. the parameter captured by the *C* variable) also plays a role. So whether Ben hits the target does *not* depend, in the relevant sense, on Ben's skills.

5.3 Truth conditions for ability ascriptions

At this point, I can give an informal statement of the truth conditions of ability ascriptions.

> ⌜*S* is able to A⌝ is true just in case:
>
> (i) there is a circumstantial world *w* such that *S* As in *w*;
>
> (ii) whether *S* As depends on *S*'s intrinsic properties, given physical laws and background facts;
>
> (iii) *w* witnesses this dependence.

Clause (i) states the at-issue truth-conditions of ability ascriptions. Clause (ii) and (iii) state a not-at-issue element. Clauses (i) and (ii) are familiar. Clause (iii) introduces something new. By saying that *w* 'witnesses' this dependence, I mean that it is included in the dependence domain. So, while the circumstantial domain and the dependence domain may fail to overlap for some ability ascriptions, they invariably overlap whenever an ability ascription is true or false (i.e., not undefined).

To see the rationale behind the overlap requirement, consider yet another variant of the dart scenario.

> **Pre-empted ability.** Ava is a skilled dart player who rarely misses a target. She's about to throw and, normally, she would be very likely to hit. But her magician friend Hermione wants to be completely sure that she doesn't miss. So, right before Ava starts to throw, she will take control of Ava's body and move her arm to produce a perfect shot, landing on target.

Several speakers judge that (32) is not true in this scenario:[17]

(32) Ava is able to hit the target on this throw.

[17] Interestingly, not all speakers share this judgement; some speakers find (32) acceptable even in this context. I conjecture that the extreme context-dependence of ability ascriptions (see §5.4) is playing a role here.

This even though (i) there is a circumstantially accessible world where Ava hits the target (in fact, she hits the target in all such worlds) and (ii) whether Ava hits the target depends, on her intrinsic properties throughout a set of worlds where we hold fixed Ava's intrinsic features and physical laws, but we exclude occasional help by magicians. What is missing? The problem, I suggest, is that this dependence is not exerted anywhere in circumstantial worlds. Clause (iii) addresses this.

The diagrams below represent the relations between the circumstantial domain and the dependence domain in various cases. Overlap is required whenever ⌜*S* is able to A⌝ is true or false. We have undefinedness whenever the dependence domain contains both A and non-A-worlds (case 1) and there is no overlap between dependence domain and circumstantial worlds (case 2).

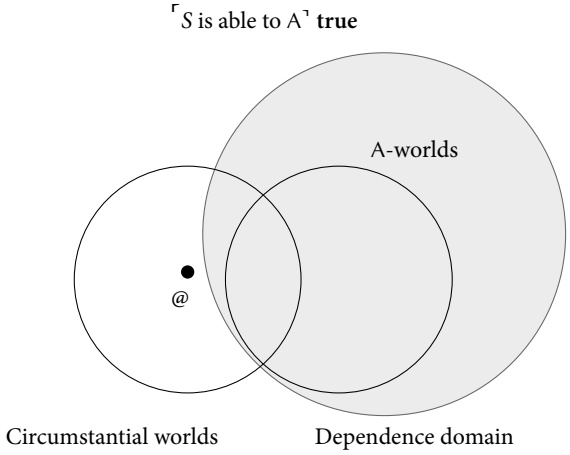

⌜*S* is able to A⌝ **false**

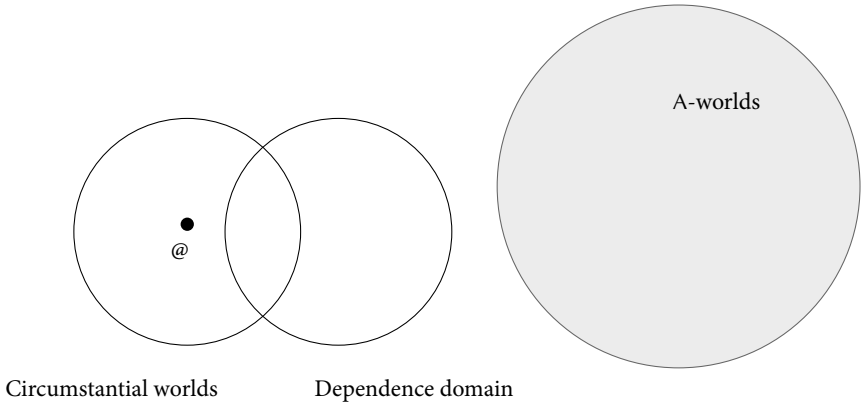

⌜*S* is able to A⌝ **undefined** (case 1)

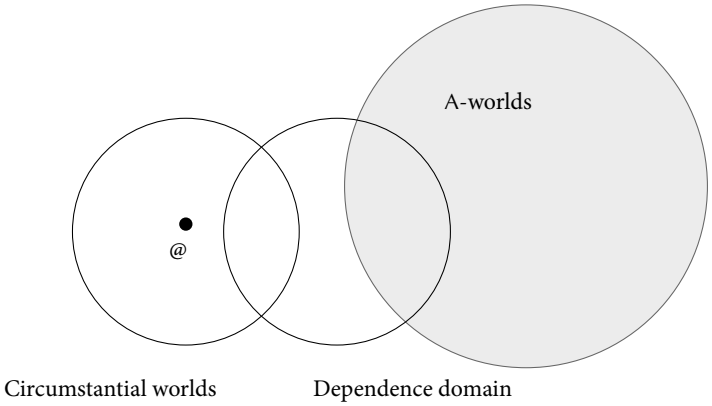

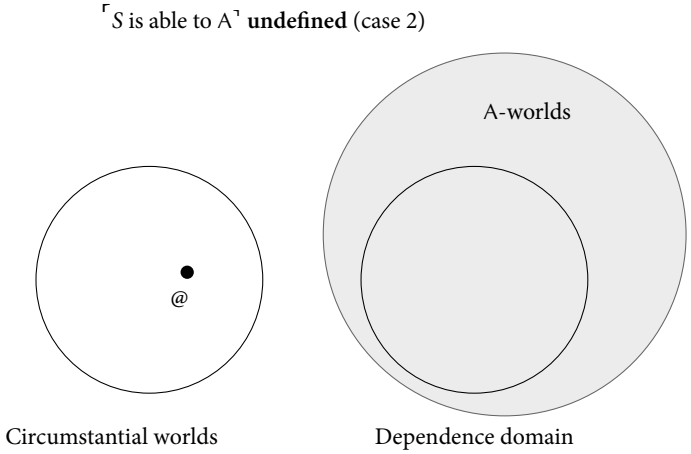

Figure 2. Relations between circumstantially accessible worlds and the dependence domain for cases of true, false, and undefined ability ascriptions.

5.4 The role of context-dependence

Before moving on to the formal semantics, let me briefly discuss context-dependence.

To define sufficiency and dependence, I have used a notion of 'background circumstances'. What are 'background circumstances'? I propose that what counts as a background circumstance is just determined by context. Hence the same ability report might be true or not depending on how this parameter is filled in by context.

This might seem a flaw of the theory. It predicts that ability reports are extremely context-dependent—even more than standard modal claims. But, upon reflection, this prediction seems exactly correct. For an illustration, suppose that Ava is considering purchasing a ticket for a fair lottery. The following is very awkward:

(28) Ava is able to win the lottery.

But now, suppose that only individuals whose first name starts with 'A' are allowed to win. (The lottery still involves a random draw between these individuals.) In this scenario, (28) can be judged true. Given that nothing in the Ava's internal states and causal powers has changed, this shift must be due to a contextual parameter being filled in differently. I suggest that, in the new context, the proposition that Ava's ticket is drawn is built into the background circumstances. So we get that there is a feature of Ava's (i.e. her first

name starting with 'A')[18] that, given favorable background circumstances (which include her ticket being drawn), is sufficient for her winning the lottery. So the sufficiency requirement is satisfied and (28) is true.

6 Formal semantics

In this section, I implement the account of §5 in a compositional semantics. As I argued in §3, ability modals involve a not-at-issue element in their meaning. For an ascription of the form ⌜*S* is able to A⌝, the content of this not-at-issue element is, roughly, that whether *S* As depends on intrinsic features of *S*, given laws and background facts. More specifically, and assuming the notion of dependence domain defined in §5, the not-at-issue element imposes a disjunctive condition:

(i) either, in all worlds in the dependence domain, *S* As;

(ii) or, in all worlds in the dependence domain, *S* does not A;

The disjunctive form of the presupposition is crucial for accounting for the extra strength of *able* in both positive and negative contexts.

I have spoken, in general terms, of a 'not-at-issue element' in the meaning of ability modals. Why not call it a presupposition? The short answer is that there are several kinds of not-at-issue meaning, and it's far from clear that this is a case of presupposition. As mentioned in §3, presuppositions project along systematic patterns in complex sentence. It is unclear that the definedness requirement of ability modals matches this projection pattern. For example, presuppositions standardly project out of conditional antecedents, as shown in (29). But the definedness requirement of *able* seemingly doesn't, as shown in (30).

(29) a. If Clio stopped smoking, then... $\rightsquigarrow$ Clio used to smoke
b. I have no idea whether Clio ever smoked. ?? But if Clio stopped smoking, her wife is happy.

(30) a. If Ava is able to hit the target now, then... $\not\rightsquigarrow$ Whether Ava hits depends on her skills

[18] From a metaphysician's perspective, of course, being named 'Ava' is not an intrinsic feature of Ava's. But, given that there are scenarios where (28) is clearly judged true, it seems to be something that speakers are willing to count it as an intrinsic property when assessing dependencies.

b. I have no idea about Ava's dart-throwing skills. But if Ava is able to hit the target now, she will win the game.

One alternative option is that the not-at-issue element is a form of homogeneity. Homogeneity is a widespread phenomenon in language, roughly concerning the interaction between some quantificational operators and negation (Križ 2015, Križ & Spector 2021, Bar-Lev 2021). A number of modal items have been linked to it.[19] Crucially for us, homogeneity conforms to the pattern of projection under negation individuated in §3, and hence a homogeneity analysis is compatible with the data of this paper.[20]

For my purposes, I don't need to settle the nature of this not-at-issue content. It is sufficiently clear that *able* has a not-at-issue component in its meaning, whatever that component turns out to be. I will treat this component as a definedness requirement on the semantic value of ability report, in line with what happens both with presupposition and with homogeneity.

6.1 The meaning of *able*

Basic framework. My starting point is a Kratzer-style framework. On Kratzer's semantics, the interpretation of modals is relativized to two conversational background parameters, a modal base, represented as 'f', and an ordering source, represented as 'g'. Both are functions from worlds to sets of propositions. The modal base specifies a set of assumptions that have to be validated by all worlds in the domain of quantification of the modal. The ordering source specifies a set of priorities, which induce an ordering on the worlds in the domain. I assume that, for ability modals, f and g are interpreted as follows.

(i) $f(w)$ is the set of propositions that are true in circumstantial worlds (hence $\bigcap f(w)$ is the set of circumstantially accessible worlds at w).

[19] See e.g. Schlenker 2004 for a homogeneity-based account of conditionals, and Agha & Jeretič 2022 for a homogeneity-based account of weak necessity modals. A *caveat*: recent accounts of homogeneity have actually moved away from the idea that homogeneity is a presupposition (see references in the text). I stick to the presuppositional analysis for convenience here; I'm very open to the idea that eventually the extra element of meaning of ability modals should be analyzed as another kind of not-at-issue content.

[20] Unfortunately, the literature on homogeneity is still developing, and we don't have yet a clear and accepted picture of homogeneity projection in various environments. See Križ 2015 for a discussion of the projection behavior of homogeneity and how it differs from presupposition.

(ii) $g(w)$ is the set of laws of nature in w. The information that was captured via structural equations in causal models (§5.2) will be encoded in propositions in this set.

Extra conversational background. I add a third conversational background parameter, which I represent as 'b'. b is also a function from worlds to sets of propositions. Intuitively, b maps a world to a set of background assumptions, which hold throughout the dependence domain (§5). I assume that the set of these background assumptions is always consistent, i.e.:

$$\text{For all } w\text{: } \bigcap b(w) \neq \varnothing$$

As usual, the ordering source is used to single out a set of 'best' worlds by inducing a ranking (Kratzer 1981, 1986, 2012). Differently from standard semantics, the domain that is ordered is not the modal base, but $\bigcap b(w)$, i.e. the set of worlds where background propositions (call it 'background set') are true. The resulting set of most highly ranked worlds $\textsc{best}_{w,b,g}$ is what I called 'dependence domain'.

Truth conditions.

(31) $[\![S \text{ is able to A}]\!]^{w,f,g,b}$ = true iff
$\exists w' \in \bigcap f(w)$ such that $[\![S \text{ As}]\!]^{w',f,g,b} = 1$:
$w' \in \textsc{best}_{w,b,g}$ and
(i) either $\forall w'' \in \textsc{best}_{w,b,g}$, $[\![S \text{ As}]\!]^{w'',f,g,b} = 1$
(ii) or $\forall w'' \in \textsc{best}_{w,b,g}$, $[\![S \text{ As}]\!]^{w'',f,g,b} = 0$

As anticipated, the truth conditions proper involve simply existential quantification over circumstantial worlds. The definedness condition states that the world that witnesses this existential quantification is in the dependence domain, and that either all worlds in the dependence domain make the prejacent true, or none does. (Technical note: the definedness condition is stated as a requirement on the variable bound by the existential quantifier in the truth conditions.)

Let's see an example in detail. Consider again:

(32) Ava is able to hit the target on this throw.

Let the conversational backgrounds at w yield:

$$f(w) = \left\{ \text{ Ava is skilled } \right\}$$

$$g(w) = \left\{ \begin{array}{l} \text{Ava throws well iff she's skilled} \\ \text{Ava hits the target iff: she throws well and there is no wind} \end{array} \right\}$$

$$b(w) = \left\{ \begin{array}{l} \text{Ava is skilled} \\ \text{There is no wind} \end{array} \right\}$$

For simplicity, suppose we have only four worlds:

w_1 : Ava is skilled, no wind, Ava hits
w_2 : Ava is skilled, there is wind, Ava hits
w_3 : Ava is skilled, no wind, Ava misses
w_4 : Ava is skilled, there is wind, Ava misses

We have the following:

(i) All of $w_1 - w_4$ are in $\bigcap f(w)$;

(ii) $\bigcap b(w) = \{w_1, w_3\}$ (since in w_2 and w_4 there is wind, which is ruled out by $b(w)$);

(iii) $\textsc{best}_{w,b,g} = \{w_1\}$ (since w_1, but not w_2, verifies all propositions in $g(w)$).

So, in this toy case, the dependence domain is just the singleton $\{w_1\}$. Since w_1 is a circumstantially accessible world where Ava hits the target, w_1 is in the dependence domain, and all worlds in the dependence domain are worlds where Ava hits, (32) is both defined and true in this case.

6.2 Predictions

Since the facts discussed in §3 we made concern inferences, we need a notion of consequence. The most natural notion we can use requires that, at all points of evaluation where the premises are defined and true, the conclusion is also defined and true.[21] Using the notion of an index as a shorthand for the n-tuple of parameters to which we relativize interpretation, we define this formally as follows:

[21]Notice that this notion is different from the notion of Strawson-entailment (see von Fintel 1999), which is widely adopted in the literature on presupposition. Strawson entailment requires that, at all points at which the premises are true *and the conclusion is defined*, the conclusion is also true. This notion of entailment allows that an inference can be valid, while there are scenarios where the premises are true and the conclusion undefined. The notion I define, conversely, is closer to preservation of truth in Kleene-style trivalent logic.

$A_1, \ldots, A_n \vDash C$ iff for all indices i such that $[\![A_1]\!]^i, \ldots, [\![A_n]\!]^i$ are defined and true, $[\![C]\!]^i$ is defined and true.

Now we can check predictions. Recall that the central gauntlet of §3 was that ability modals appear stronger than circumstantials in both positive and negative contexts.

Ability asymmetrically entails circumstantial possibility.

$\diamondsuit\!\!\!\cdot A \vDash \Diamond A$	*S* is able to A $\vDash$ *S* can A
$\Diamond A \nvDash \diamondsuit\!\!\!\cdot A$	*S* can A $\nvDash$ *S* is able to A

Inability asymmetrically entails circumstantial impossibility.

$\neg \diamondsuit\!\!\!\cdot A \vDash \neg \Diamond A$	*S* is not able to A $\vDash$ *S* cannot A
$\neg \Diamond A \nvDash \neg \diamondsuit\!\!\!\cdot A$	*S* cannot A $\nvDash$ *S* is not able to A

The entailment direction, in both cases, is straightforward. The truth-conditional part the meaning of *S is able to A* is simply that the prejacent is circumstantially possible. The entailments from ability to circumstantial possibility, and from inability to circumstantial impossibility, immediately follow from this.

Now let's check that we get failure of entailment from circumstantial possibility to ability ($\Diamond A \nvDash \diamondsuit\!\!\!\cdot A$). For illustration, consider once more the case of Ben, the unskilled dart thrower.

(33) a. It can happen that Ben hits the target on this throw.
b. Ben is able to to hit the target on this throw.

Take a very ordinary context, where Ben hits the target in some circumstantial worlds and misses in others. (33a) is true. At the same time, intrinsic features of Ben's do not determine whether he hits. So the definedness requirement is not satisfied, and the sentence is undefined.

Finally, we get failure of entailment from circumstantial impossibility to inability ($\neg \Diamond A \nvDash \neg \diamondsuit\!\!\!\cdot A$). Consider again:

(34) a. The baby cannot fall.
b. The baby is not able to fall.

Take a context where the carrier is indeed very safe, and hence there is no circumstantial world where the baby falls. (34a) is true. Assume that, differently from the modal base, the background parameter b does not encode

information about the safety of the carrier. So some worlds in the dependence domain are safe-carrier worlds, and some are not. Given plausible assumptions about dependencies, the baby will fall in some of these worlds and not in others. So the definedness requirement for (34b) is unsatisfied, and the sentence is undefined.

7 Refuting the agentive analysis

The dependence analysis contrast with the currently dominant line of theorizing about ability, which ties ability to agency. In this section, I sketch the agentive analysis and show that it has serious shortcomings. Just the cases that motivated the dependence analysis play a major role.

The agentive idea is pithily put by Malte Willer (2021):

> 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 . . . In fact, there is good reason to think that the role of agency is key to generating our [logical] puzzles . . . 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. (Willer 2021, p. 555)

Without going into details, it's useful to state, schematically, agentive truth conditions for ability reports. Virtually all agentive analyses yield a semantics that conforms to the following template, where '⇒' denotes a relation between propositions:

(35) $[\![S \text{ is able to A}]\!]$ is true iff there is a suitable action a such that:

$$S \text{ performs } a \;\Rightarrow\; [\![S \text{ As}]\!] \text{ is true}$$

Agentive analyses mainly differ in how they cash out the '⇒' operator. The literature includes two main options. Conditional analyses (e.g. Mandelkern, Schultheis, and Boylan 2017, henceforth MSB) take it to be a Stalnaker-style conditional relation: the prejacent of the modal is required to be true in the closest world where the agent performs the action. Double modal analyses (e.g. Brown 1988, Fusco 2021, Willer 2021) take it to be a necessitation relation: the performance of the action necessitates, on some relevant modal flavor, the prejacent.

The agentive analysis is arguably intuitive, and connects to several features of ability modals. In particular, it provides a route for explaining why

able is stronger than circumstantial *can*. On the agentive analysis, *S is able to A* conveys not merely that it's circumstantially possible that *S* As, but also that there is an action that *S* can perform that necessitates that *S* As. Some traditional and contemporary accounts of the extra strength of *able* (e.g. Brown 1988, Fusco 2021, and Willer 2021) build just on this idea.[22]

Despite these promising features, and despite the wide endorsement of agentive analyses, I claim that the semantics of ability modals involves no reference to agency. I give three arguments.

7.1 Problem #1: circumstantial necessity doesn't entail ability

Recall the Ben case from §4:

> **Magical dart.** Ben is a mediocre dart thrower who's about to throw a dart. In ordinary circumstances, there would be a high chance that he would miss. But Ben's magician friend Camille wants Ben's dart to hit the target. So, as soon as the dart leaves Ben's hands, Camille will cast a spell on the dart, leading it to the target.

We judge that (18b) is true, but (19) is not:

(18b) Ben cannot miss the target on this throw.

(19) # Ben is able to hit the target on this throw.

This shows that not only circumstantial possibility, but also circumstantial necessity fails to entail ability:

$$\Box A \not\models \diamondsuit A$$

The first problem for the agentive analysis is that, *modulo* a plausible assumption, it is committed to this entailment. Hence it predicts truth conditions that are too weak, and overgenerates.

The 'plausible assumption' is, roughly, that the domain of quantification of '$\Rightarrow$' is restricted to circumstantially accessible worlds.[23] Unfortunately, agentive theorists often leave the flavor of the relevant modality implicit. But, on the assumption that 'suitable actions' have to be circumstantially accessible, this seems very plausible.

[22] Historically, the agential analysis has been deployed to capture Kenny's puzzle (1976), which is related to the strength puzzle. See §8 for discussion of Kenny's puzzle.

[23] More precisely, the assumption is: $\Box_{\text{Circ}}(A \supset B)$ entails $A \Rightarrow B$,

Let me consider one particular account, i.e. Mandelkern, Schultheis, and Boylan's (2017). The MSB account is admirably clear: '⇒' is a Stalnaker-style counterfactual. Hence an ability claim is true just in case, for some suitable action a[24], a Stalnaker counterfactual of the form of (36) is true:

(36) If Ben performed *a*, he would hit the target.

There are no established principles linking circumstantial and counterfactual domains. But, on the assumption that suitable actions are circumstantially possible, it seems obvious that (36) should quantify over circumstantially accessible worlds. In that case, whenever (18b) is true, (36) is predicted to be true too. Moreover, conditionals of the form of (36) *do* sound true in the Ben scenario (after all, Ben will hit no matter what!). So MSB's account wrongly validates the inference from circumstantial necessity to ability.

7.2 Problem #2: the negative part of the strength puzzle

The second argument builds on the second example from §4. Recall:

> **The baby carrier.** Someone is carrying their infant daughter in a baby carrier. The baby is leaning out in a way that appears dangerous, and you worry that she might fall. But the carrier is actually very safe.

(37) is true, while (38) is defective.

(37) The baby cannot fall.

(38) # The baby is not able to fall.

These judgments illustrate the failure of the entailment from $\neg \Diamond A$ to $\neg \diamondsuit A$. Agentive analysis, again on plausible assumptions, cannot capture this.

Here are the schematic truth conditions for (38), on the agentive analysis:

(39) $[\![(38)]\!]$ = true iff there is no suitable action *a* such that

The baby performs *a* ⇒ The baby falls

With some approximation, these truth conditions can be glossed as: *The baby won't fall, no matter what she does*. Again, the key issue is what possibilities '⇒' quantifies over. If '⇒' is restricted to circumstantial possibilities,

[24]MSB define a notion of a 'practically available action'. The details are unimportant here.

the agentive analysis wrongly predicts an entailment between (37) and (38). As I argued above, this restriction seems plausible.

There is also a further difficulty for this case. Independently of the interaction between circumstantial and ability modals, it seems impossible to get a true and felicitous reading for (38). (38) just sounds awkward.[25] This is unexpected on the agentive analysis: if (38) means that the baby won't fall in any case in which she performs any number of actions, it should be perfectly acceptable and, in the right contexts, it should even be true. This is a further respect in which the agentive analysis overgenerates.

7.3 Abilities without agents or actions

The last argument is the simplest one. *Able* routinely appears in sentences where the subject is not an agent, and where the relevant verb phrase does not describe an action. Below are some examples.

Ability without agents. Some ability reports simply have subjects that are not agents.[26]

(40) a. Steel is able to withstand a pressure of 100 tons.
b. This type of coal is able to burn without producing smoke.
c. The central mall on campus is able to hold 10,000 people.

Ability without actions and control. The second batch of cases show that ability reports can, and routinely do, involve complements that denote events that are not actions. Here are some examples.

(41) a. Ava is able to fall asleep quickly this evening.
b. Ben is able to fight off the virus over the next days.
c. Clem is able to process sugar very fast now (she's on a special medication).

[25] According to some speakers, (38) has a marginal reading that is felicitous. Suppose that the baby is trying to disentangle herself from the carrier, in the attempt of throwing herself on the ground. Then, according to these speakers, (38) can be heard as felicitous. But my point is that (38), but not (37) or the claim that the baby won't fall no matter what, is infelicitous in the more ordinary context where the baby is quietly sitting in the carrier.

[26] Thanks to *[name omitted]* for first pointing out to me examples of this sort. Some speakers find the data in (40) jarring, but Google searches show that these readings are genuine. A search for "steel is able to" returns 976,000 hits, one for "wood is able to" returns 1,790,000 hits.

d. Dave is able to hear unusually high frequencies now (he's on drugs).
e. Emma is able to digest poisonous mushrooms today (she took an antidote).

Falling asleep, fighting off a virus, and processing sugar are not actions. They are not the result of a deliberative process and their occurrence doesn't require that an agent should be aware that they are happening. Yet (41a)–(41e) are perfectly natural ability reports.

These examples also cast doubt on a further claim. It is often said that, in some sense or other, ability requires control (see e.g. Loets & Zakkou 2022 for recent discussion of this). This claim may be made precise in many ways, depending on what notion of control one has in mind. But (41a)–(41e) are counterexamples to the claim that ability requires control, on *any* plausible notion of control. This is not surprising, once we divorce the notion of ability from the notion of agency.

8 Further logical puzzles

In closing, I discuss some other logical features of ability modals, to show that the dependence analysis captures them.

8.1 Failure of Distribution over Disjunction

Kenny 1976 famously presents a version of this case:

> **Context.** Clem is a somewhat experienced dart player. She's good enough to reliably hit the board when she throws, but she is still doesn't control what section of the board she hits.

Kenny observes that (42a) sounds true, but (42b) and (42c) don't.

(42) a. Clem is able to hit the top or the bottom part of the board on this throw.
b. Clem is able to hit the top part of the board on this throw.
c. Clem is able to hit the bottom part of the board on this throw.

But the following is a standard principle of modal logic, which is also validated by all standard modal semantics:

Distribution over Disjunction (DoD). $\Diamond(A \vee B) \vDash \Diamond A \vee \Diamond B$

But now, via **DoD**, (42a) should entail the disjunction of (42b) and (42c). The fact that the former is true and the latter two are false suggests that **DoD** also fails for ability modals.

The dependence analysis correctly predicts failures of **DoD**. The key point is that the following can obtain: whether Clem hits the board depends on Clem, but whether she hits the top part of the board (or the bottom part of the board) does not.[27] When this happens, the definedness condition of (42a) is satisfied, but the definedness conditions of each of (42b) and (42c) are not. So the inference is correctly predicted to be invalid.

8.2 Success inferences

Consider now the following inference pattern:[28]

Success. $A \vDash \diamondsuit A$ $\qquad\qquad$ *S* As $\vDash$ *S* is able to A

Most theorists hold that **Success** fails. To see this, suppose that Ben, the mediocre dart player, is about to throw. A quick peak at a crystal ball shows us that, this time, Ben will get very lucky: he's going to hit the target.[29] Still, even in this context we are not going to assent to the following:[30]

(43) # Ben is able to to hit the target on this throw.

The failure of **Success** dovetails with another logical fact, which is reminiscent of the asymmetries noticed in §3: the contrapositive of **Success** is valid (as noticed in Boylan 2022). To see this, notice that conjunctions like (44) sound contradictory:

(44) # Ava is not able to hit the target on this throw, but she will hit the target on this throw (by sheer luck).

[27] In terms of the formal analysis spelled out in §6, what needs to happen is the following: all worlds in the dependence domain are worlds where Clem hits the board; but some of them are worlds where she hits the top part, and half of them are worlds where she hits the bottom part.

[28] Failures of success were first noticed by Kenny 1976. For extensive discussion of **Success**, see Boylan 2022.

[29] It doesn't matter whether the crystal ball is actually accurate, but only that we believe it is.

[30] For reasons of space, I have to keep my argument for the failure of **Success** very concise. For an extended and convincing argument, see Boylan 2022. In that paper, Boylan claims that **Success** holds for past ability claims. Arguably, this is related to the phenomenon of actuality entailments, which is fully orthogonal to the semantics of ability modals (see, among many, Hacquard 2020). I must leave a full discussion of Boylan's proposal to future work.

In sum, we have a situation mirroring our findings in §3:

Actuality does not entail ability.	$\mathsf{A} \not\models \diamondsuit\mathsf{A}$	*S* As $\not\models$ *S* is able to A
Inability entails inactuality.	$\neg\diamondsuit\mathsf{A} \models \neg\mathsf{A}$	*S* is not able to A $\models$ *S* won't A

Also in this case, the dependence analysis makes the correct predictions. **Success** fails, because it has counterexamples whenever *S* As, but *S*'s A-ing does not depend on *S*. Conversely, whenever *S is not able to A* is true and defined, *S As* is false in all circumstantial worlds. Since the latter include the actual world, *S doesn't A* is also defined and true.

9 Conclusion

I have defended an account of ability modals based on a simple idea: *able* and related expressions track dependencies between an event ad the subject of ability. I have shown that this idea can be made precise and can be implemented in a compositional semantics that is surprisingly conservative.

The dependence idea runs against a long tradition of thinking of ability modals as related to agency. As I have argued, this is exactly as it should be. Agency analyses get many things right, but they are insufficiently general and don't yield a full account of the strength puzzle. The dependence analysis, conversely, handles all cases with ease.

References

Agha, Omar, and Paloma Jeretič (2022). "Weak necessity modals as homogeneous pluralities of worlds." In *Proceedings of Semantics and Linguistic Theory*, vol. 32, p. 831–851. Linguistic Society of America.

Arregui, Ana (2004). *On the accessibility of possible worlds: The role of tense and aspect*. Ph.D. thesis, University of Massachusetts at Amherst.

Arregui, Ana (2007). "When Aspect Matters: The Case of Would-Conditionals." *Natural Language Semantics*, 15(3): pp. 221–264.

Arregui, Ana (2009). "On Similarity in Counterfactuals." *Linguistics and Philosophy*, 32(3): pp. 245–278.

Bar-Lev, Moshe E. (2021). "An Implicature Account of Homogeneity and Non-Maximality." *Linguistics and Philosophy*, 44(5): pp. 1045–1097.

Bird, Alexander (2007). *Nature's Metaphysics: Laws and Properties*. Oxford University Press.

Boylan, David (2022). "Does Success Entail Ability?" *Noûs*, 56(3): pp. 570–601.

Brown, Mark A. (1988). "On the Logic of Ability." *Journal of Philosophical Logic*, 17(1): pp. 1–26.

Cariani, Fabrizio (2021). *The Modal Future: A Theory of Future-Directed Thought and Talk*. Cambridge, UK: Cambridge University Press.

Cariani, Fabrizio, and Simon Goldstein (2018). "Conditional Heresies." *Philosophy and Phenomenological Research*, (2): pp. 251–282.

Cariani, Fabrizio, and Paolo Santorio (2018). "Will Done Better: Selection Semantics, Future Credence, and Indeterminacy." *Mind*, 127(505): pp. 129–165.

Carnie, Andrew (2021). *Syntax: A Generative Introduction (Fourth Edition)*. Wiley-Blackwell.

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

Fusco, Melissa (2021). "Agential Free Choice." *Journal of Philosophical Logic*, 50(1): pp. 57–87.

Hackl, Martin (1998). "On the Semantics of "Ability Attributions"." Unpublished draft available at `http://web.mit.edu/hackl/www/papers/files/AbilityAttributions.pdf`, MIT.

Hacquard, Valentine (2006). *Aspects of modality*. Ph.D. thesis, Massachusetts Institute of Technology.

Hacquard, Valentine (2010). "On the Event Relativity of Modal Auxiliaries." *Natural Language Semantics*, 18(1): pp. 79–114.

Hacquard, Valentine (2020). "Actuality Entailments." In *The Wiley Blackwell Companion to Semantics*, Wiley Blackwell.

Halpern, Joseph (2000). "Axiomatizing Causal Reasoning." *Journal of Artificial Intelligence Research*, 12: pp. 317–337.

Halpern, Joseph (2016). *Actual Causality*. MIT Press.

Heim, Irene (1983). "On the projection problem for presuppositions." In *Proceedings of WCCFL*, vol. 2, pp. 114–125. Stanford: Stanford University Press.

Herzberger, Hans G. (1973). "Dimensions of Truth." *Journal of Philosophical Logic*, 2(4): pp. 535–556.

Karttunen, Lauri (1974). "Presupposition and linguistic context." *Theoretical Linguistics*, 1(1–3): pp. 181–194.

Kaufmann, Stefan (2013). "Causal Premise Semantics." *Cognitive Science*, 37(6): pp. 1136–1170.

Kenny, Anthony (1976). "Human abilities and dynamic modalities." In *Essays on explanation and understanding*, Dordrecht: Reidel, doi=https://doi.org/10.1007/978-94-010-1823-4$_1$1, pp. 209–232.

Klecha, Peter (2013). "Diagnosing modality in predictive expressions." *Journal of Semantics*, 31(3): pp. 443–455.

Kratzer, Angelika (1977). "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy*, 1(3): pp. 337–355.

Kratzer, Angelika (1981). "The Notional Category of Modality." In H. J. Eikmeyer, and H. Rieser (eds.) *Words, Worlds, and Contexts: New Approaches to Word Semantics*, Berlin: de Gruyter.

Kratzer, Angelika (1986). "Conditionals." In *Chicago Linguistics Society: Papers from the Parasession on Pragmatics and Grammatical Theory*, vol. 22, pp. 1–15. University of Chicago, Chicago IL: Chicago Linguistic Society.

Kratzer, Angelika (2012). *Modals and Conditionals: New and Revised Perspectives*, vol. 36. Oxford University Press.

Kratzer, Angelika (2013). "Modality for the 21st Century." In *The Language Cognition Interface / L'Interface Language-Cognition*, Geneva-Paris (Librairie Droz), pp. 179–199.

Križ, Manuel (2015). *Aspects of homogeneity in the semantics of natural language*. Ph.D. thesis, University of Vienna.

Križ, Manuel, and Benjamin Spector (2021). "Interpreting Plural Predication: Homogeneity and Non-Maximality." *Linguistics and Philosophy*, 44(5): pp. 1131–1178.

Leslie, Sarah-Jane, and Adam Lerner (2022). "Generic Generalizations." In E. N. Zalta, and U. Nodelman (eds.) *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University. Fall 2022 ed.

Lewis, David K. (1997). "Finkish Dispositions." *Philosophical Quarterly*, 47(187): pp. 143–158.

Loets, Annina, and Julia Zakkou (2022). "Agentive Duality Reconsidered." *Philosophical Studies*, 179(12): pp. 3771–3789.

Mackie, J.L. (1965). "Causes and Conditions." *American Philosophical Quarterly*, 2(4): pp. 245––264.

Mandelkern, Matthew (2023). *Bounds: Meaning and the Limits of Interpretation*. Forthcoming with *Oxford University Press*.

Mandelkern, Matthew, Ginger Schultheis, and David Boylan (2017). "Agentive Modals." *Philosophical Review*, 126(3): pp. 301–343.

Nadathur, Prerna (2021). "Causality and aspect in ability, actuality, and implicativity." In *Proceedings of Semantics and Linguistic Theory*, vol. 30, p. 409–429. Linguistic Society of America.

Pearl, Judea (2000). *Causality: models, reasoning and inference*. Cambridge University Press.

Santorio, Paolo (2019). "Interventions in Premise Semantics." *Philosophers' Imprint*, 19.

Schaffer, Jonathan (2016). "Grounding in the Image of Causation." *Philosophical Studies*, 173(1): pp. 49–100.

Schlenker, Philippe (2004). "Conditionals as definite descriptions." *Research on language and computation*, 2(3): pp. 417–462.

Schlenker, Philippe (2008). "Be articulate: a pragmatic theory of presupposition projection." *Theoretical Linguistics*, 34(3): pp. 157–212.

Schwarz, Wolfgang (2020). "Ability and Possibility." *Philosophers' Imprint*, 20.

Spencer, Jack (2017). "Able to Do the Impossible." *Mind*, 126(502): pp. 466–497.

Sudo, Yasutada (2012). *On the semantics of phi features on pronouns*. Ph.D. thesis, Massachusetts Institute of Technology.

Vetter, Barbara (2015). *Potentiality: From Dispositions to Modality*. Oxford, England and New York, NY, USA: Oxford University Press.

Willer, Malte (2021). "Two Puzzles About Ability Can." *Linguistics and Philosophy*, 44(3): pp. 551–586.

Wilson, Alastair (2018). "Metaphysical Causation." *Noûs*, 52(4): pp. 723–751.