

Ability and Possibility*

Wolfgang Schwarz

Published in *Philosophers' Imprint* 20(6): 2020

Abstract. According to the classical quantificational analysis of modals, an agent has the ability to perform an act iff (roughly) relevant facts about the agent and her environment are compatible with her performing the act. The analysis faces a number of problems, many of which can be traced to the fact that it takes even accidental performance of an act as proof of the relevant ability. I argue that ability statements are systematically ambiguous: on one reading, accidental performance really is enough; on another, more is required. The stronger notion of ability plays a central role in normative contexts. Both readings, I argue, can be captured within the classical quantificational framework, provided we allow conversational context to impose restrictions not just on the “accessible worlds” (the facts that are held fixed), but also on what counts as a performance of the relevant act among these worlds.

1 Introduction

It has long been noticed that modals in natural language behave much like quantifiers over possible worlds.¹ In this tradition, David Lewis [1976: 150f., 1979: 246f.] and Angelika Kratzer [1977] influentially suggested that ‘*S* can ϕ ’ (or ‘*S* is able to ϕ ’) is true just in case there are accessible worlds at which *S* ϕ s, where a world is *accessible* just in case it is compatible with contextually relevant facts about the actual world.

The quantificational analysis works well for statements like (1), from [Kratzer 1991].

- (1) Hydrangeas can grow here.

In a suitable context, an utterance of (1) conveys that relevant facts about the regional climate, soil, etc. are compatible with the growth of hydrangeas. In other words, the utterance is true iff hydrangeas grow at some possible worlds where these facts are held fixed.

* Earlier versions of this paper were presented at a workshop in Berlin and departmental seminars in Aberdeen, Canberra, and Glasgow. I thank the audiences at these events, as well as Mathias Böhm, Stephanie Collins, Fausto Carcassi, Romy Jaster, Theodore Korzukhin, Brian Rabern, Michael Ridge, Clas Weber, and two anonymous referees for helpful comments.

¹ See [Copeland 2002] for a partial history.

However, when we turn to ‘can’ statements attributing abilities to intentional agents, the quantificational analysis appears to run into trouble. Suppose Maisy wants to open a safe, but doesn’t know the combination.² Intuitively, (2) is then false.

(2) Maisy can open the safe.

On the quantificational analysis, this means that there are no accessible worlds at which Maisy opens the safe. But all it takes to open the safe is to dial the right combination. Let’s say the combination is 448-961-5237. Surely Maisy can dial that combination. No relevant facts about her inner state or her environment, it seems, are incompatible with her dialling 448-961-5237. But if there are accessible worlds where Maisy dials 448-961-5237, how could there fail to be accessible worlds where she opens the safe?

Examples are easy to multiply. Suppose Cyril does not know the first 10 digits of π . Intuitively, (3) is then false.

(3) Cyril can recite the first 10 digits of π .

But, we may assume, Cyril can utter any permutation of the numerals ‘three’, ‘one’, ‘four’, etc. One of these permutations lists the the first 10 digits of π . So it looks like there *are* accessible worlds at which Cyril recites the first 10 digits of π .

One might try to get around these problems by tinkering with the accessibility relation. We will see that this isn’t easy. It also doesn’t seem to get at the heart of the problem. Intuitively, when we say that someone can recite the first 10 digits of π , we don’t just mean that no relevant facts preclude them from uttering ‘three, one, four,’ etc. Rather, the agent must have a certain kind of intentional control over performing the act under the description of ‘reciting digits of π ’. G.E. Moore’s [1912] conditional analysis gets closer to what is required. According to Moore, ‘*S* can ϕ ’ is true just in case *S* would ϕ if *S* tried to ϕ . If Cyril tried to recite the first 10 digits of π , he would almost certainly fail. Similarly, Maisy would almost certainly fail if she tried to open the safe.

To be sure, the conditional analysis also faces some well-known problems. Among other things, it looks unappealing for sentences like (1) or (4) or (5), which don’t seem to imply the kind of intentional control implied by (2) and (3).

(4) My phone can connect to the WiFi network.

(5) Steel can withstand temperatures of 1000 °C.

Some have concluded that modal expressions like ‘can’ and ‘able’ are ambiguous between an *agentive* sense and a non-agentive or *circumstantial* sense. On this view, the classical quantificational analysis may be adequate for the circumstantial sense, but not for the agentive sense. Many attempts have been made to analyse the agentive sense,

² This example is borrowed from [Carlson 1999].

but no consensus has emerged from these efforts.³ We still lack a clear picture of what ‘can’ means in (2) and (3), and how this meaning is related to the meaning in (1), (4), and (5).

I will argue that despite initial appearance, (2) and (3) can be analysed in terms of simple, existential quantification over accessible worlds. (2) is false because there are no accessible worlds at which Maisy opens the safe – provided that accidental openings don’t count. In general, I will argue that when we evaluate whether someone can or cannot ϕ , conversational context may put restrictions on what counts as a relevant way of ϕ ing. Sometimes accidental ϕ ing counts, sometimes not. As a result, ability statements are ambiguous between a weaker reading, on which accidental ϕ ing suffices, and a logically stronger reading, on which non-accidental ϕ ing is required.

That there is some such ambiguity is, I think, independently plausible. For example, although (2) is most naturally interpreted as false in the described scenario, there is also a sense in which it is true. Maisy *can* open the safe, insofar as she can dial 448-961-5237, and that’s what she needs to do in order to open the safe. The false reading is more salient, in part because the true reading is trivial: we would normally take for granted that Maisy has enough control over her arms and fingers to dial whatever combination happens to be the right one.⁴

The ambiguity is easier to see in cases where neither reading is trivial. Suppose Eddie’s favourite piano piece is the Moonlight Sonata; Charley can play the Moonlight Sonata, but he doesn’t know that it is Eddie favourite piece. Now consider (6).

(6) Charley can play Eddie’s favourite piece.

True or false? Intuitively, the statement is true in one sense and false in another. Informally, Charley can play what is in fact Eddie’s favourite piece, but he can’t play it “under that description”. In contrast to (2), here the weak reading is far from trivial: most people can turn the dial on a combination lock, but few can play the Moonlight Sonata.⁵

³ See, among others, [Cross 1986], [Brown 1988], [Peacocke 1999], [Mele 2003], [Vihvelin 2004], [Fara 2008], [Maier 2015], [Maier 2018b], [Mandelkern et al. 2017], [Jaster 2020].

⁴ In addition, when we consider a sentence like (2), we are often interested in what the agent ought to do. As we will see, this makes the strong reading salient.

⁵ In (6), the description ‘Eddie’s favourite piece’ can arguably take either narrow scope or wide scope over the modal ‘can’. If the description is given wide scope, as enforced in (6’), the sentence is intuitively true.

(6’) Eddie’s favourite piece is something Charley can play.

One might conjecture that the ambiguity of ability modals is generally a *de re/de dicto* ambiguity. However, (2) does not contain a relevant description (‘the safe’ doesn’t help), yet it intuitively displays the same ambiguity as (6). Here it’s as if the entire verb phrase ‘open the safe’ can scope over the modal, so that (2) asserts *of opening the safe* that Maisy can do *it*. I will return to the relevance of scope in footnote 29 below.

At the other end of the spectrum lie statements like (7).

(7) Usain Bolt can run 100 meters in 9.58 seconds.

Here the weak reading is more salient. If someone utters (7), they plausibly mean that Usain Bolt can do something that amounts to running 100 meters in 9.58 seconds – a highly non-trivial feat. The claim is not that Bolt is able to perform the act “under that description”, meaning that he can at will choose to run with an average speed of exactly 0.0958 m/s.

It will be useful to have labels for the two readings. If an agent has the ability to ϕ in the weaker sense, I will say that she can ϕ *effectively*. Maisy, for example, can do something – namely, dial 448-961-5237 – that *in effect* amounts to opening the safe, although Maisy is not aware of the connection between dialling 448-961-5237 and opening the safe. If an agent has the ability to ϕ in the stronger sense, I will say that she can ϕ *transparently*.⁶

I take both of these readings to express an agentive type of modality. Compare (7) with (8), said about a statue precariously placed on a ledge.

(8) The statue can easily fall.

The truth of (8) has nothing to do with the statue’s agency; (8) simply states that a certain kind of accident could easily befall the statue. By contrast, (7) does not say that a certain kind of running event could easily befall Usain Bolt. If Bolt runs a race, he is exercising his agency.

Indeed, (2), (3), (6), and (7) could all be paraphrased with ‘able’ (unlike (8)), and the same ambiguity would arise. The ambiguity also affects other ability modals like ‘in a position to’, as well as corresponding necessity modals like ‘must’. It also arises for deontic modals. Suppose Maisy is required by some normative code to dial 448-961-5237, and (by coincidence) dialling 448-961-5237 is tantamount to opening the safe. In one sense, then, ‘Maisy ought to open the safe’ is true, although Maisy is not required to perform the act “under that description”.

I hope you see the ambiguity in (2), (3), and (6). Let’s try to get clearer about the two readings. What is required for an agent to have an ability in the weak, effective

⁶ [Mandelkern et al. 2017: 320f.] also notice the effective/transparent ambiguity for a sentence like (2). They suggest that the two readings are analogues of the objective/subjective ambiguity of deontic modals, arising from a traditionally neglected information sensitivity of ability modals. However, we will see in example (9) below that possession of a transparent ability can depend on external matters of which both agent and attributer are unaware. Moreover, while the subjective reading of deontic modals is sensitive to information supplied by ‘if’-clause (as highlighted in [Kolodny and MacFarlane 2010]), the reading on which (2) is false is not: on its most salient reading, (2’) is just as false as (2).

(2’) If the combination is 448-961-5237, then Maisy can open the safe.

Third, as I’m just about to point out, the effective/transparent ambiguity arguably also affects deontic modals, where it is orthogonal to the subjective/objective ambiguity.

sense? And what is required to have an ability in the strong, transparent sense? In the next section, I will sketch a preliminary answer, tailored to a certain use of ability statements in normative contexts. In the remainder of the paper, I will explore how the two readings might fit into the classical quantificational paradigm. To this end, I will first review the quantificational treatment of modals in section 3, and defend it against a certain misunderstanding. In section 4 I will motivate the idea that the interpretation of modals is sensitive not just to contextual restrictions on the accessible worlds, but also to restrictions on the interpretation of the embedded verb phrase (the “prejacent”). The transparent reading of ability statements is then easily explained by a non-accidentality restriction. In section 5 I will sketch a few possibilities for how the non-accidentality restriction may be compositionally derived. Finally, in section 6, I will show that the proposed account defuses several prominent objections and counterexamples to the quantificational analysis of ability modals.

2 Oughts and cans

To get a better grip on the effective/transparent ambiguity, it will help to begin with a particular class of ability statements, in which the subject is an intentional agent facing a particular choice.

In such a case, it is natural to ask not just what the agent *can* do but also what she *ought* to do, or what she is *obligated* to do.⁷ The two questions are related. If an agent isn’t able to ϕ , we typically wouldn’t judge that she ought to ϕ : ‘ought’ implies ‘can’, as the saying goes. Moreover, ‘ought’ plausibly implies ‘can’t do better’: an agent ought to ϕ only if there is no alternative act she could perform that is better than ϕ .⁸

So normative judgements about what an agent ought to do are sensitive to the what the agent can do, in some sense of ‘can’. This gives us a job description for an analysis of ‘can’: the acts an agent can perform should be the options that are candidates for what she ought to do. Let’s see what ‘can’ should mean if it is to fit that job description.

Here is a simple picture, or model, of rational agency. Voluntary acts, we assume, are consequences of corresponding intentions or volitions. When an agent deliberates about what to do and finally reaches a decision, the direct result is not an overt act, but a volitional state – an intention, or an all-things-considered desire, to perform a certain act. Which volitional state comes about depends on the agent’s desires, values, beliefs, and the process of deliberation. Other values, other beliefs, or other deliberational processes might lead to other volitional states and consequently to other acts.

⁷I will ignore the difference between ‘ought’ and ‘obligated’ in what follows.

⁸[Portmore 2019: ch.1] argues that any credible moral theory supports this connection between ‘ought’ and ‘can’t do better’, possibly with minor caveats.

Normative judgements about what an agent ought to do are generally relative to norms and facts: in light of such-and-such norms and such-and-such facts, the agent ought to ϕ . The judgement does not presuppose that the agent knows the relevant facts or that she adheres to the relevant norms. Thus it can easily happen that an agent ought to ϕ (relative to some norms and facts) and yet chooses some alternative ϕ' because relative to her own information or her own values, ϕ' is preferable to ϕ . Roughly speaking, when we ask what an agent ought to do (relative to such-and-such norms and facts), we want to know what she *would* do if she were adequately responsive to the relevant reasons. Any act the agent could have performed if only she were differently motivated should therefore count as an available option. An act is disqualified as an option only if no (reasonable) variation of the agent's volitional state would lead her to perform the act.

The parenthetical reasonability condition is meant to exclude volitional states that are in a certain sense “unavailable” to the agent. If an agent has severe arachnophobia, we may not want to say that she ought to pick up a spider, even if doing so would have desirable consequences. Here the volitional state of intending to pick up the spider is treated as unavailable. I will have a little more to say on the role of the availability restriction, but I will not try to spell out exactly what it involves.⁹

Let's say that an act is *under an agent's volitional control* just in case there is an available variation of the agent's volitional state that would cause her to perform the act. The picture of agency and normativity just outlined suggests that when we consider what an agent ought to do, the relevant options – what the agent can do, in the normatively relevant sense of ‘can’ – are precisely the acts that are under the agent's volitional control.

But things are more complicated. Return to Maisy, standing in front of the safe. On the present proposal, Maisy *can* open the safe (in the normatively relevant sense), for there is an available variation of her volitional state – namely, intending to dial 448-961-5237 – that would cause her to open the safe.

But suppose Maisy could prevent some tragedy by opening the safe. One might have thought that if an agent can perform an act that would prevent a tragedy (at no significant costs), and if no other available act would be comparably good, then the agent is obligated to perform that act. More cautiously, one might have thought that at least if an agent *knows* that she can perform an act that would prevent a tragedy etc., then she is obligated to perform the act. But on the proposed analysis of ‘can’, Maisy *can* prevent the tragedy by opening the safe, and she *knows* that she can. (She knows, we

⁹ There are difficult normative questions here. Just how strongly must you be afraid of spiders so that picking up the spider no longer counts as an available option? Can ordinary intentions make incompatible intentions unavailable, as in Goldman's [1978] example of an agent who supposedly ought to accelerate given her ill-advised intention to change lanes? Some have argued that there are different notions of ‘ought’, corresponding to different rules for what is treated as available (e.g. [Timmerman and Cohen 2016]). If that is right, we might want to allow for correspondingly different notions of ‘can’. The account I will develop is neutral on these questions.

may assume, that some available variation of her volitional state would cause her to open the safe.) Yet it seems wrong to claim that Maisy is obligated to open the safe. We certainly wouldn't blame her if she doesn't open the safe. So when we consider what Maisy ought to do, we do not seem to count all the acts that are under her volitional control as relevant options.¹⁰

In the previous section, I observed that sentences like (2) are ambiguous. Maisy *can* open the safe in the weaker “effective” sense, but she *can't* open the safe in the more demanding “transparent” sense. Plausibly, if an act is under an agent's volitional control, then the agent can perform it in the effective sense. To fit the normative job description, we need something stronger. This is, I think, where the transparent concept of ability comes into play.

As a first stab, a transparent ability to ϕ requires that the agent can ϕ simply *by deciding to ϕ* . Since she doesn't know the combination, Maisy can't open the safe simply by deciding to open the safe. While there is an available volition that would cause her to open the safe – namely, intending to dial 448-961-5237 – Maisy doesn't know which of the available volitions would do the job. Similarly for Cyril in (3) and Charlie in (6). What Maisy, Cyril, and Charley lack is neither skill nor opportunity, but plain old information. Due to their lack of information, they don't know what they would have to do in order to perform the desired act.

So here is a tentative and preliminary analysis for the two readings of ‘*S can ϕ* ’ (and ‘*S is able to ϕ* ’). The analysis will be generalised in section 4.

Analysis 1

- a. *S can ϕ* (effectively) iff there is an available variation of *S*'s volitional state that would cause *S* to ϕ .
- b. *S can ϕ* (transparently) iff there is an available variation *V* of *S*'s volitional state such that (i) *V* would cause *S* to ϕ , and (ii) *S* knows that *V* would cause her to ϕ provided that ϕ ing is under her volitional control.

I will explain the ‘provided that’ clause in a moment. First I want to explain how abilities in the sense of Analysis 1b fit the normative job description of delineating the options that are candidates for what an agent ought to do.¹¹

¹⁰ One might bite the bullet and say that Maisy is indeed obligated to open the safe, even though she isn't blameworthy for flouting the obligation. As I will argue below, this is not only counter-intuitive, it also fails to respect the “action-guiding” or “advisory” role of oughts and obligations. If you'd nonetheless prefer to bite the bullet, much of what I'm going to say could be recast in terms of a connection between ‘can’ and blame rather than a connection between ‘can’ and ‘ought’.

¹¹ If ‘ought’ implies ‘can’, and the relevant sense of ‘can’ involves condition (ii) of Analysis 1b, we

Consider what is required for an agent to be under an obligation to perform a certain act. Arguably, the following two conditions are necessary (though not sufficient). First, performing the act must be under the agent’s volitional control: some possible variation of her motivational state would make her perform the act. That’s what condition (i) gives us. But more is required. In addition, the agent must know what she would have to do in order to perform the act. Condition (ii) ensures that if an agent can ϕ in the relevant sense, then she knows what she has to do in order to ϕ (ignoring the proviso).

Exactly why that second condition is required depends on one’s views about normativity, but here is a plausible sketch of an explanation. A central role of normative judgements is often glossed as “motivational”, “action-guiding”, or “advisory”. There is a close connection between the assumption that an agent ought to ϕ and the assumption that it would be appropriate to advise the agent to ϕ . But if an agent has no idea what she would have to do in order to ϕ , then it would not be appropriate to tell her to ϕ . If you (as an advisor) know the safe’s combination, you might reasonably advise Maisy to dial 448-961-5237. Perhaps that is what Maisy ought to do in the “objective” sense of ‘ought’ – it is what she ought to do in light of all the facts. But advising Maisy to open the safe would not be reasonable, no matter how good it would be if she opened the safe. An advice only makes sense if the recipient has some idea how she could follow the advice.

The advisory function of normative judgements also explains why there should be an availability restriction on the relevant motivational variations. If an agent has severe arachnophobia, there is little point in advising her to pick up a spider. When giving advice, physiological conditions that are insensitive to reason must be accepted as constraints.

Now let me explain the ‘provided that’ clause in Analysis 1b. Consider the following scenario. Dotty has been led into a room with a single door. There is nobody else in the room, the door is unlocked, and Dotty could freely leave. However, she falsely believes (for good reasons, if you want) that the door is locked. Can she open the door? Intuitively, yes. That is, (9) does not fit into the pattern of (2), (3), and (6).

(9) Dotty can open the door.

In contrast to the earlier examples, there is no relevant reading on which (9) is false. Dotty *believes* that she can’t open the door, but her belief is simply false. Dotty can open the door, and she can do it simply by deciding to open the door.

Yet Dotty does *not* know that deciding to open the door would cause her to open the door. On the contrary, she believes that nothing she could decide would cause her to

might expect to find a trace of that condition in the semantics of ‘ought’. Indeed, recent accounts of deontic modals have appealed to some such condition in order to give an adequate treatment of the “miner puzzle” of [Kolodny and MacFarlane 2010]. For example, [Charlow 2013] suggests that the semantics of deontic modals involves an operation that filters out “non-actionable” propositions, where a proposition ϕ is actionable iff there is some “available act” A such that the relevant information entails $A \supset \phi$. (See also [Kolodny and MacFarlane 2010: 132, fn.28].)

open the door. Without the proviso in clause (ii), Analysis 1b would predict that (9) is false (on the transparent reading of ‘can’); it would falsely predict that (9) patterns with (2), (3), and (6).

The proviso gets around the problem. Dotty does not know that deciding to open the door would cause her to open the door, but she plausibly does know that *if opening the door is under her volitional control, then* deciding to open the door would cause her to open the door. Informally, if she were informed that she can open the door, she would know how to go about achieving that result.

Notice how the judgement about (9) depends on the assumption that there is only one door. Suppose there are three doors; Dotty believes that all of them are locked, but in fact one is unlocked. For some reason, she can’t simply try all the doors to get out: if she were to try a locked door, an alarm would ring and the unlocked door would lock as well. With these changes, Dotty no longer knows what she would have to do in order to leave the room, even on the supposition that it is in her power to leave the room. Analysis 1b predicts (correctly, I think) that Dotty no longer has the ability to leave the room in the strong, transparent sense. She can leave the room, but only by luck.

How does the proviso in Analysis 1b fit the normative job description? Again, the details depend on one’s views about normativity, but superficially the proviso goes well with the advisory picture. It would not be reasonable to advise Maisy to open the safe, but it would be perfectly reasonable to advise Dotty (in the original scenario) to open the door.¹²

Analysis 1b is closely related to Moore’s conditional analysis, so a brief comparison may be in order.¹³ Suppose an agent has an ability to ϕ in the sense of Analysis 1b. That is, some variation V of her volitional state would cause her to ϕ , and she knows that V would have this effect (provided that ϕ ing is under her volitional control). Under normal circumstances, trying to ϕ would then cause her to ϕ ; so the agent also has the ability to ϕ in the sense of the conditional analysis. An exception might be cases where the agent knows that she can perform a given act only by trying to perform a different act. I am not sure whether, in such a case, the agent intuitively has or lacks the relevant ability (in the transparent sense), and whether the act should count as an option for normative purposes. As it stands, Analysis 1b says the agent has the ability, while the conditional analysis says she does not.

The two accounts more clearly come apart in the other direction. It can easily happen that an agent would ϕ if she tried to ϕ without having the ability to ϕ in the sense

¹² Of course, the advice would only be appropriate from a perspective in which it is known that the door is unlocked. Thus in a more “subjective” sense, relative to Dotty’s own information, it may still not be true that Dotty ought to open the door.

¹³ Like Analysis 1b, Moore’s analysis arguably targets the normatively relevant sense of ‘can’. It is not a coincidence that it is presented in a book called *Ethics*.

of Analysis 1b. Notoriously, the conditional analysis implies that arachnophobiacs can pick up spiders, or that coma patients can get up and leave their bed (see [van Inwagen 1983: 119]); Analysis 1b does not have these consequences. More tellingly perhaps, the conditional analysis faces problems with accidental success. Suppose if Maisy were to attempt to open the safe, 448-961-5237 would be the first combination she would try – by sheer luck. Would that render (2) unambiguously true? Arguably not. Or suppose Maisy in fact tries to open the safe and happens to dial 448-961-5237. Assuming that ‘*A* and *B*’ entails ‘if *A* were the case then *B* would be the case’ (a popular, although not entirely uncontested assumption in the logic of counterfactuals), the conditional analysis implies that (2) is unambiguously true. But arguably, it is not.¹⁴

Some have suggested that just as one can try-and-succeed without having the relevant ability, one can try-and-fail even though one has the ability.¹⁵ John Austin’s classical example is of a golfer who kicks himself for missing a putt, knowing that he had the ability to hole it (under the given circumstances). That kind of case seems possible to me. Analysis 1b gets it right, provided the golfer knows what he would have had to do in order to make the putt.¹⁶

Note that in these cases, the conditional analysis not only fails to match certain linguistic intuitions. More importantly, it fails to satisfy the normative job description of identifying the acts an agent might be obligated to perform. Maisy is not obligated to open the safe even if she would succeed, by chance, if she tried.

By comparison, consider another apparent problem for the conditional analysis: that it does not account for attributions of so-called “general abilities”, as in (10).

(10) Charley can play the piano.

Suppose Charley, a skilled piano player, is currently on an Arctic expedition on which he does not have access to a piano. There’s a natural sense in which (10) is true, although it is not the case that Charley would play the piano if only he tried. This is not a problem if our target is the normatively relevant sense of ‘can’ (and ‘able’) that delineates an agent’s options in a given choice situation. Without access to a piano, Charley can’t be obligated to play the piano; while he is on the expedition, playing the piano is not one of his options.

What (10) illustrates is that neither Analysis 1 nor the conditional analysis will do as a fully general analysis of ability statements. We’ve known this all along. A general

¹⁴ The judgement that even successful ϕ ing does not entail an ability to ϕ – at least in some normatively important sense of ‘ability’ – is common in the literature; see e.g. [Kenny 1976: 214], [Fara 2008], [Vihvelin 2013: 182], [Southwood and Wiens 2016], [Maier 2018b]. Nonetheless, it sounds odd to say that an agent is unable to ϕ and yet ϕ s; I will return to this observation in section 5.

¹⁵ See [Austin 1961], [Fara 2008], [Vihvelin 2004: 187], [Mandelkern et al. 2017: 312].

¹⁶ Must the golfer be able to identify the volitional state that would have made him succeed? See page 18.

analysis would also have to cover statements like (1), (4), and (5), which intuitively have nothing to do with tryings or volitions.¹⁷

3 The quantificational canon

Let me recap where we stand. In the previous section, we looked at ability statements in which the subject is an intentional agent facing a particular choice. Such statements have a logically stronger and a weaker reading. I have proposed an analysis of the two readings, which, I argued, fits an important role of ability statements in normative contexts. We have also seen that the proposed analysis does not cover ability statements like (1), (4), (5), or (10), which do not concern particular choices of intentional agents.

A tempting response is to postulate further ambiguities. In the philosophical literature, it is common to distinguish between attributions of “specific abilities”, concerning an agent’s options in a given choice situation, and attributions of “general abilities” like (10) that do not imply a present opportunity to exercise the ability.¹⁸ So perhaps ability modals have not two, but four meanings: the effective and transparent reading of “specific” ability modals captured by Analysis 1, a third meaning for “general” ability statements like (10), and a fourth (“circumstantial”?) meaning for (1), (4), and (5).

On closer inspection, however, the effective/transparent ambiguity also arises for “general” and “circumstantial” ability statements. Take (3): ‘Cyril can recite the first 10 digits of π ’. This is naturally understood as an attribution of a “general” ability: the statement may be true even if Cyril does not presently have an opportunity to recite the digits. But on that interpretation, the statement is intuitively still ambiguous between a (relatively trivial) effective reading and a (non-trivial) transparent reading. So we seem to have an at least six-fold ambiguity.

Things get worse if we try to spell out what some of the other meanings might be. Return to (10). When we say that someone can play the piano, we typically mean that they have certain lasting, intrinsic properties characteristic of piano players. In that sense, (10) can be true even if Charley’s arm is broken so that he is temporarily unable to exercise his ability. In another sense, however, a broken arm renders (10) false. In yet another sense, (10) can be true even if Charley has not taken any piano lessons, as long as he could acquire the capacity to play the piano with sufficient training and practice. (Imagine Charley has a disability, and we consider which musical instruments

¹⁷ John Maier [2015, 2018a] suggests that ability statements of all kinds can be analysed in terms of an agent’s options. This is an overgeneralisation, I think. I am also not sure how Maier’s concept of an option, which he takes as primitive, is related to the normative job description I have outlined. For example, Maier [2018a: 424] states that whenever an agent ϕ s, even by luck, then ϕ ing was an option for the agent. This suggests that Maier’s options do not stand in any simple relation to judgements about what the agent ought to do.

¹⁸ See e.g. [Austin 1961: 230], [Honoré 1964], [Whittle 2010], [Mandelkern et al. 2017], [Maier 2018a].

his disability would allow him to play.) Alternatively, one could use (10) to convey that Charley not only has the capacity to play the piano, but could also bring about an opportunity with reasonable effort: returning from his expedition, Charley (still in the airport) might say how glad he is that he can now once again play the piano.

So there isn't a single "general" reading of (10). Depending on context, there seems to be an almost open-ended range of things one can express with a sentence like (10).¹⁹

Yet other uses of 'can' and 'able' don't seem to fit into any of the above six categories. Consider (11) and (12).

(11) I can't come to the party; I need some rest.

(12) I can't come to the party; I have to look after my children.

These are "specific" ability statements insofar as they concern a particular choice of an intentional agent, but their most salient interpretation is not captured by Analysis 1 (or the conditional analysis).

At this point, the strategy of postulating brute ambiguities begins to look unappealing. It is hard to believe that words like 'can' and 'able' just happen to have dozens of distinct meanings. (And we haven't even touched on deontic or epistemic uses of 'can'.) At the very least, we would like to know what all these meanings have in common.

Here the quantificational analysis of modals shows its strength. On the quantificational analysis, 'can' and 'able' belong to a larger family of modal constructions all of which are analysed in terms of quantification over a restricted domain of ("accessible") possible worlds. An agent *can* ϕ iff she ϕ s at some accessible worlds; an agent *must* ϕ iff she ϕ s at all accessible worlds; and so on. The analysis is meant to cover all uses of 'can', from (1) to (12).

Let's be clear about what the quantificational analysis is trying to achieve. Take a particular ability, such as Charley's ability to play the Moonlight Sonata. What makes it the case that Charley has this ability? Intuitively, the ability is grounded in certain physiological features relating to finger dexterity, hand coordination, muscle memory, etc. – features Charley has acquired through years of piano practice. The quantificational account instead seems to suggest that Charley's ability is grounded in facts about other possible worlds: Charley has the ability to play the Moonlight Sonata, here in our world, because there is some other world where he (or worse, someone else, see [Lewis 1986: ch.4]) plays the Moonlight Sonata. That sounds wrong.²⁰

Similar complaints have been raised against other applications of possible-worlds semantics, such as the Lewis-Stalnaker account of counterfactuals, or the analysis of physical necessity in terms of nomically possible worlds.²¹ But these complaints misunderstand

19 Essentially the same point is made in [Kratzer 1981].

20 See e.g. [Vetter 2013] for this kind of complaint.

21 See also the "Humphrey objection" in [Kripke 1980: 44f.].

the aim of possible-worlds semantics. Charitably understood, the analysis of modal constructions in terms of possible worlds is not meant to identify the metaphysical grounds of the relevant phenomena. Nor is it meant to offer a substantive reduction of the phenomena to psychologically more basic concepts. It is merely supposed to offer a useful framework for modelling the phenomena themselves. As Lewis explains, in response to the objection that counterfactuals are made true by facts about our world:

[It] is indeed the character of our world that makes the counterfactual true. But it is only by bringing the other worlds into the story that we can say in any concise way what character it takes to make what counterfactuals true. The other worlds provide a frame of reference whereby we can characterise our world. [Lewis 1986: 22]

In the quantificational analysis of ‘can’ and ‘able’, the substance lies in the accessibility relation. Typically, accessibility is a matter of preserving certain features of the agent and her environment. When we say that Charley can play the Moonlight Sonata, for example, we hold fixed Charley’s finger dexterity, muscle memory, etc. It is Charley’s possession of these features, here in our world, that explains why he plays the Moonlight Sonata in some accessible world and therefore why he can play the Moonlight Sonata back in our world. As in the case of counterfactuals, the other worlds merely provide a convenient frame of reference for characterising the actual world.

What is convenient about the quantificational analysis becomes apparent once we look at more examples and a wider range of contexts. Recall the many things one could convey by (10). When we consider whether someone can play the piano, we sometimes hold fixed a broken arm; sometimes we don’t. Sometimes we hold fixed lack of training and practice; sometimes we don’t. Sometimes we hold fixed circumstantial absence of a piano; sometimes we don’t. Metaphysically, there is little all these readings of (10) have in common, and there is even less they have in common with (5), (9), or (11). By refraining to give a substantive analysis, the quantificational analysis is able to bring out a common theme.

The simple quantificational analysis that I have presented has undergone many refinements in recent decades. For example, Kratzer [1981, 1991] argued that the worlds over which modals quantify are sometimes restricted not just by relevant facts but also by the satisfaction of salient norms or goals. This shows up in (11) and (12). When we evaluate ‘I can’t come to the party’, we plausibly hold fixed the speaker’s location, her obligations, her needs, etc. But that is not enough to render (11) or (12) true. After all, there are worlds where I come to the party despite my contrary needs or obligations. One could postulate that these worlds are not accessible. But then accessibility could no longer be understood in terms of facts we intuitively hold fixed. Kratzer instead makes the interpretation of ‘can’ relative to a second contextual parameter: an “ordering source”

which ranks possible worlds by the extent to which they satisfy relevant norms or goals or standards. ‘ S can ϕ ’ is true (roughly) iff S ϕ s in some of the top-ranked worlds among those compatible with the facts we hold fixed.

Glossing over some further refinements, the lexical entry for ‘can’ may therefore look approximately as follows:²²

$$[[\text{can}]]^{w,R,\leq} = \lambda P_{\langle s, \langle e, t \rangle \rangle} \lambda x_e. \exists w' (wRw' \wedge \forall w'' (wRw'' \rightarrow w'' \leq w') \wedge P(w)(x) = 1).$$

Here, w is the world of evaluation, R is a contextually supplied relation encoding the facts about w that are held fixed, and \leq is a contextually supplied partial order that ranks w' above w'' just in case w'' better satisfies relevant norms or goals or standards. (If no standards are salient, all worlds are ranked equally.) ‘ S can ϕ ’ is true at w iff S ϕ s in some of the highest-ranked worlds R -related to w .

The quantificational analysis not only brings out a common theme in ability statements; it also explains how the ability sense of ‘can’ is related to the ‘can’ of permission (‘can I leave now?’), the epistemic ‘can’ (‘you can’t have seen her, she’s overseas’), and genuinely circumstantial or “participant-external” uses of ‘can’ in statements like (8) or (13) that can’t be paraphrased with ‘able’ (in contrast to all other examples in this paper).

(13) It can take years to earn someone’s trust.

On the quantificational account, all of these involve existential quantification over possible worlds. Necessity modals like ‘must’ are similarly analysed in terms of universal rather than existential quantification. Thus we can explain why, say, ‘I must sneeze’ is equivalent to ‘it is not the case that I can not sneeze’.

The quantificational account has other attractive features. For example, it allows an attractive explanation for the interaction between ‘can’ and ‘if’ in sentences like ‘I can dance if nobody is watching’, where, on the most natural interpretation, the if-clause has neither wide scope nor narrow scope, but rather seems to restrict the set of worlds over which the modal quantifies (see [Kratzer 1986], compare [Horgan 1979: 350]). The account also explains why ‘ought’ implies (a specific kind of) ‘can’: if S ϕ s in the normatively best of the contextually accessible worlds – as ‘ S ought to ϕ ’ states – it follows that S ϕ s in some contextually accessible worlds – as ‘ S can ϕ ’ states.

In sum, the popularity of the quantificational analysis in formal semantics is not a historical accident. If we wanted to trade the quantificational analysis of (2) for, say, a

²²I here assume that ‘can’ is of type $\langle\langle s, et \rangle, et \rangle$, taking as argument a verb phrase that denotes a property; Kratzer instead assumes a “raising” analysis on which ‘can’ is of type $\langle st, t \rangle$. The raising analysis helps to account for “participant-external” uses of ‘can’ like (8) and (13) below, where ‘can’ scopes above tense and aspect. Genuine ability modals, however, are always “participant-internal”, scoping below tense and aspect. Whether separate lexical entries are required for participant-external and internal modals is an open question. See [Brennan 1993] and [Hacquard 2010] for relevant discussion.

broadly conditional analysis (as many recent accounts of agentive modals suggest), we would need an explanation of what this conditional sense of ‘can’ has to do with the obviously related sense of ‘can’ in (7), (10), (11), and (12), let alone (1), (5), and (8). We would also need a new explanation of the connections between ‘can’, ‘must’, and ‘ought’, of the interactions between ‘can’ and ‘if’, and a long list of other facts I haven’t even mentioned.

On the other hand, the quantificational analysis of ability modals seems to face a number of decisive objections. One of these is the problem I mentioned in the introduction (I will turn to others in section 6): the analysis does not seem to capture the transparent reading of sentences like (2), (3), and (6). Let me go through the problem again.

Maisy is standing in front of a safe, whose combination – unbeknown to Maisy – is 448-961-5237. Intuitively, (14a) is true while (14b) is false (on its more salient transparent reading).

- (14) a. Maisy can dial 448-961-5237.
- b. Maisy can open the safe.

The quantificational analysis easily accounts for the truth of (14a): since Maisy has the skill and opportunity (although possibly not the motivation) to dial 448-961-5237, there are accessible worlds where she dials that combination. But then how could there fail to be accessible worlds where Maisy opens the safe, so that (14b) comes out false? We would have to assume that when we evaluate (14a), worlds where Maisy opens the safe by dialling 448-961-5237 are accessible, but these same worlds become inaccessible when we evaluate (14b). But why? What other worlds become inaccessible when we evaluate (14b)? Remember that accessibility is supposed to track certain facts about Mary that we intuitively hold fixed. What are these facts in the case of (14b), and why can’t we hold them fixed when we evaluate (14a), which does not have a salient false reading? Why does (14b) become true if Maisy gains information about the safe’s combination?

As we saw in section 2, when we look at Maisy’s options, it is natural to hold fixed Maisy’s general physiology and macroscopic environment, because these are insensitive to reason, while allowing her volitional state to vary. But that would render (14b) true. Indeed, whatever facts we hold fixed, the actual world itself will always come out as accessible (from itself). Yet in section 2 we saw that on the normatively relevant reading, ‘*S* can ϕ ’ may be false even if *S* actually ϕ s. So we would have to invoke an ordering source. Perhaps worlds at which Maisy performs an act by luck or accident are ranked lower because these acts don’t conform to the ideal standard of intentional agency? The problem is that any world at which Maisy intentionally dials 448-961-5237 is a world at which she unintentionally opens the safe. So invoking the suggested ordering source would predict that (14a) is false as well.

This is not an impossibility proof. All I claim is that I don’t see a natural way to

account for the transparent reading of ability modals by tinkering with the accessibility relation or other parameters in the standard quantificational model. We need to introduce another parameter.

4 Restricting the prejacent

When we talk about whether someone can or cannot ϕ , we sometimes have in mind a particular way of ϕ ing. Consider a well-known puzzle from action theory. Tallulah’s left arm is paralysed. So (15) is intuitively false.

(15) Tallulah can move her left arm.

But suppose Tallulah’s right arm is not paralysed. As a consequence, she *can* move her left arm, the way she can move a cup on the table: by grabbing it with her right arm. That’s not the kind of moving we have in mind when we deny (15). When we judge that Tallulah can’t move her left arm, we mean that she can’t move her left arm “actively”, using the muscles in the arm. On its most salient interpretation, the truth of (15) therefore seems to require not just that there are accessible worlds at which Tallulah moves her left arm. There must be worlds at which she moves the arm “actively”. Holding fixed Tallulah’s paralysis, there are no such worlds. Hence (15) is false.

Another puzzle that points in the same direction concerns the quality threshold for ability attributions. Assume it took Charley many years of practice to become good at the piano. At what point along this journey did (10) (‘Charley can play the piano’) become true? Even before his first lesson, Charley was able to produce sounds by pushing the keys on a piano. But this is usually not what we have in mind when we say that someone can (or cannot) play the piano. We mean that they can (or cannot) play *reasonably well*. What that means depends on conversational context. Simplifying a little, let’s imagine that every instance of piano playing can be assigned a scalar performance level taking into account both the difficulty of the piece and the quality of the play.²³ We might then explain the context-dependence of (10) by assuming that context fixes a minimal performance level so that (10) counts as true iff there are accessible worlds at which Charley plays the piano at a level that exceeds the threshold.

So there is evidence that when we evaluate ‘ S can ϕ ’, conversational context can restrict what counts as a relevant act of ϕ ing among the accessible worlds. This, I suggest, is also how the transparent reading of ability statements comes about.

Return once more to Maisy in front of her safe. Since Maisy can dial 448-961-5237, by ordinary standards there will be accessible worlds at which she opens the safe. But in all these worlds she opens the safe by luck. When we judge that Maisy can’t open the safe,

²³ The simplification is that these two aspects, and various sub-aspects in which they could be decomposed, can be measured on a single scale. I also ignore the obvious vagueness in (10).

we don't mean that she can't do it by luck. We mean that she can't open the safe in the deliberate, controlled, conscious, non-accidental way in which she could open it if she knew the combination.

It is not terribly important for my proposal how the relevant way of opening the safe should be spelled out, but let me venture a suggestion, in line with what I suggested in section 2.

Consider what would change if Maisy knew the combination. Surveying the worlds that would result from some variation of her volitional state, we would then find not just worlds where she opens the safe by luck, but worlds where she *knowingly* opens the safe – that is, where she opens the safe while knowing that she is opening the safe.

However, imagine Maisy (not knowing the safe's combination) is wired to a brain scanner that would immediately signal to her that she is about to open the safe once it detects an intention to dial 448-961-5237. In this scenario, Maisy knowingly opens the safe in any relevant world where she forms the intention to dial 448-961-5237. Yet intuitively, the presence of the brain scanner makes no difference to the falsity of (2).

So it matters *how* Maisy comes to know that she is opening the safe in the accessible worlds. Loosely speaking, if Maisy knew the combination, she would know that she will open the safe without drawing on further information from the world.

I therefore suggest the following schematic analysis.

Analysis 2

- a. S can ϕ (effectively) iff there are (highest-ranked) accessible worlds at which $S \phi$ s.
- b. S can ϕ (transparently) iff there are (highest-ranked) accessible worlds at which $S \phi$ s *transparently*.

$S \phi$ s *transparently* iff $S \phi$ s as a result of a volitional state that warrants believing that she will ϕ provided that ϕ ing is under her volitional control.

Following Kratzer, I assume that a contextual ordering sometimes restricts the modal domain, as in (11) and (12); hence the restriction to 'highest-ranked' worlds. The 'provided that' clause in the definition of transparency mirrors the 'provided that' clause in Analysis 1. Remember (9): Dotty can open the door (transparently) even though no decision Dotty could make would, by itself, warrant believing that she will open the door. Nonetheless, a decision to (try to) open the door would warrant believing that *if she can open the door*, then she will open the door. By Analysis 2b, this is enough to render (9) true.

In contrast to Analysis 1, Analysis 2 leaves open how the accessible worlds are determined. In the kinds of cases we studied in section 2, where we consider what an

agent should do in a given choice situation, a world is accessible iff it is under the agent's volitional control – that is, iff it would result from some available variation of the agent's volitional state.²⁴ In other cases, other facts are held fixed. (3), for example, allows for a similar range of uses as (10), depending on which facts about Cyril and his environment we hold fixed. Each of these uses has an effective and a transparent reading.

Even if the standards for accessibility are determined by what's under the agent's volitional control, Analysis 2b is not fully equivalent to Analysis 1b. For one thing, Analysis 1b seems to require that the agent can somehow pick out in advance the volitional state(s) that would make her ϕ ; Analysis 2b does not. Moreover, while Analysis 1b requires knowledge, Analysis 2b only requires having a warrant for (true) belief. Among other things, the two analyses therefore give different verdicts about Gettier cases: if you're looking at a stopped clock that happens to show the correct time, and you have reason to believe the clock works, Analysis 2b says that you have the (transparent) ability to tell the time, while Analysis 1b says that you don't. In these respects, it seems to me that Analysis 2b is superior to Analysis 1b, but I won't argue the point.

As I said, I am not particularly wedded to the precise definition of transparency in Analysis 2b. My proposal is that ' S can ϕ ' is true on its transparent reading iff there are accessible worlds at which S ϕ s in a particular way – a way that plays an important role in normative contexts. For the sake of concreteness, I have tried to spell out that way, but I would not be surprised if my attempt needs revisions or amendments.

I have used an artificial adverb 'transparently' in Analysis 2b, but it may well turn out that the relevant restriction matches some more familiar concept. In particular, it may be worth exploring whether abilities in the transparent sense can be analysed as abilities to *intentionally* perform the relevant act, in some ordinary sense of 'intentionally'. According to one influential tradition in action theory, intentional action is distinguished precisely by the agent having non-observational knowledge of what she is doing (see e.g. [Anscombe 1957], [Velleman 1989], [Setiya 2008]).

Let's briefly consider what Analysis 2 says about "compulsion modals", the dual of ability modals. These are often expressed with 'must', 'have to' or 'cannot but':

(16) I have to sneeze when I look at the sun.

(17) One cannot but admire her determination.

As Mandelkern et al [2017] point out, many prominent analyses of ability modals (such as the conditional analysis) make unacceptable predictions about compulsion modals, assuming these are equivalent to 'not can not'.

Analysis 2 allows for two kinds of compulsion modals. If 'can' is read effectively, an agent cannot not ϕ iff the agent ϕ s at all (highest-ranked) accessible worlds. This seems

²⁴Here I'm assuming that there is always a single world that would determinately result from any available volitional state. This is arguably an idealisation.

to be the standard interpretation of compulsion modals. On the transparent reading, an agent cannot not ϕ iff the only (highest-ranked) accessible worlds at which the agent doesn't ϕ are worlds where she non-transparently doesn't ϕ . Since the “non-transparent” reading of compulsion modals is more common, but ability modals are often interpreted transparently, we can explain why negating an ability modal often does not entail a corresponding compulsion modal: *Maisy can't open the safe, but we may be reluctant to say that she cannot but fail to open the safe. She could get lucky.*²⁵

5 The mechanism

I have argued that when we evaluate ‘*S* can ϕ ’, conversational context can restrict what counts as an instance of ϕ ing among the accessible worlds. “Passive” arm movements are excluded in (15), inept piano performances in (10), and non-transparent safe openings in (2). But how exactly does context achieve this? What is the underlying linguistic mechanism?

Technically, we could include the restriction in the accessibility relation. There are R -accessible worlds at which an agent ϕ s ψ ly just in case there are R^* -accessible worlds at which the agent ϕ s, where R^* is defined (for example²⁶) as follows:

$$wR^*w' \Leftrightarrow wRw' \text{ and the agent } \phi\text{s } \psi\text{ly at } w'.$$

So we could in principle get the transparent reading of ‘can’ statements by applying the $*$ operation to the standard accessibility relation determined by whatever circumstances are held fixed.

This “explanation” not only robs the accessibility relation of its intuitive meaning, it also relocates the requested mechanism into the wastebasket of pragmatics. Somehow or other, it assumes, conversational context may settle that the $*$ operation is applied to the relevant facts. Note also that if R^* is the accessibility relation, then ‘*S* can ϕ ’ is true iff there are accessible worlds, full stop. We don't have to require that the agent ϕ s at some of these worlds; the definition of R^* already ensures that.²⁷ Clearly, accessibility is here doing too much work.

It would be better to add the required restriction as a separate parameter, making the interpretation of ability modals depend on the relevant circumstances, given by R ,

²⁵ Why do we normally interpret compulsion modals effectively? I'm not sure. Perhaps it's simply because the transparent reading is more complicated and – unlike transparent ability modals – does not serve any obvious purpose.

²⁶ Other definitions of R^* would also work, but none of them, I think, is much better than the one I chose.

²⁷ As a corollary, if we assume that the effective/transparent ambiguity also arises for necessity modals (as I argued on p.4), and that the same accessibility relation is often available for both types of modals (as seems required to explain various entailment facts), we get the false prediction that the transparent ‘can’ is equivalent to the transparent ‘must’.

possibly an ordering to deal with cases like (11) and (12), and a further contextual restriction C imposed on the modal domain (compare [Dorr and Hawthorne 2013]). The new parameter C would have to be highly sensitive to the modal’s prejacent: when we evaluate (14b) (‘Maisy can open the safe’), the parameter normally rules out worlds at which Maisy opens the safe by luck; when we evaluate (14a) (‘Maisy can dial 448-961-5237’), the same restriction is unavailable. Again, one would like to know more about the mechanism behind this behaviour.

An more perspicuous option would introduce a parameter ψ of adverbial type that doesn’t restrict the modal domain, but the interpretation of the prejacent. Intuitively, ψ supplies a relevant *way* of performing an act: actively, transparently, etc. The full semantics of ‘can’ might then look more or less as follows:

$$[[\text{can}]]^{w,R,\leq,\psi} = \lambda P_{\langle s, \langle e, t \rangle \rangle} \lambda x_e. \exists w' (wRw' \wedge \forall w'' (wRw'' \rightarrow w'' \leq w') \wedge \psi(P)(w)(x) = 1).$$

Semantically, the modifier ψ is simply passed on to the embedded verb phrase: S can^{actively} ϕ iff S can perform the act of actively ϕ ing.

On this approach, we don’t have to postulate that the choice of prejacent in (14a) and (14b) somehow determines a different value of the contextual parameter: if ψ imposes a transparency condition, (14a) is true and (14b) false, without any change in contextual parameters.

Other explanations might also be worth considering. For example, one might suggest that the relevant verb phrases are simply ambiguous, so that the effect has nothing to do with modals. On that approach, ‘open the safe’, for example, can mean either *open the safe transparently* or *open the safe in some way or other*. We get the transparent reading of ‘can open the safe’ if conversational context makes the first reading salient. To handle (15) and (10), we would similarly have to assume that ‘move’ is ambiguous between *move actively* and *move in some way or other*, and that ‘play the piano’ has a variety of readings with different degrees of demandingness.

An advantage of this approach is that would explain the oddity of statements like (18).

(18) Maisy wasn’t able to open the safe, but she did open the safe.

According to Analysis 2b, an agent can lack a transparent ability even if she accidentally performs the relevant act. (In section 2, I argued that this is in line with the normative job description for transparent abilities.) So (18) should sound fine. But it does not. On the present proposal, this could be explained by a preference for a uniform resolution of the ambiguity in ‘open the safe’. (18) is then ambiguous between (19a) and (19b), both of which are contradictory on the quantificational analysis (in the absence of an ordering source).²⁸

²⁸ Thanks to Brian Rabern for drawing my attention to this point.

- (19) a. Maisy wasn't able to open the safe transparently, but she did open the safe transparently.
b. Maisy wasn't able to open the safe in some way or other, but she did open the safe in some way or other.

However, the oddity of (18) is arguably related to that of (20), which the ambiguity hypothesis doesn't explain.

- (20) Maisy was able to open the safe, but she didn't open the safe.

In the semantics literature, the phenomenon illustrated by (18) and (20) is known as “actuality entailment”. [Hacquard 2009], drawing on [Bhatt 2006], offers a promising explanation within the quantificational framework that does not postulate an ambiguity of the prejacent.

In any case, one would like to see independent evidence for the supposed ambiguities. Indeed, the ambiguity hypothesis implies that (21) has a false reading if Maisy opened the safe by luck, not knowing the combination.

- (21) Maisy opened the safe.

As far as I can tell, such a reading does not exist. Even if Maisy only opened the safe by luck, (21) is unambiguously true.

Instead of postulating lexical ambiguities, one might appeal to a pragmatic or semi-pragmatic mechanism. [McCawley 1978] pointed out that the use of a simple verb phrase sometimes implicates that an event takes place in a stereotypical way. Along these lines, one might suggest that the transparent reading of ‘can ϕ ’ arises as a conversational implicature. Unlike ordinary conversational implicatures (and unlike the cases in [McCawley 1978]), however, this supposed implicature seems to arise even in questions: if Maisy doesn't know the combination, the most natural answer to ‘can Maisy open the safe?’ is ‘no’.

Another (semi-)pragmatic process that might do the job is “free enrichment”. A number of authors (including [Hall 2008], [Carston 2010], [Recanati 2013]) have suggested that when we process utterances, we often supplement the uttered sentence by unarticulated constituents that don't have to be pronounced because they can be taken for granted in the relevant context. Thus one might suggest that ‘Maisy can open the safe’ is often “enriched” to ‘Maisy can open the safe transparently’, whereas (for some reason) no corresponding enrichment is possible for ‘Maisy did open the safe’.

I have to leave a more thorough investigation of these and other possible mechanisms for another occasion. Whatever the underlying linguistic mechanism, I hope I have made it plausible that the transparent reading of agentive modals might arise through a contextual restriction on the interpretation of the embedded verb phrase – by which ‘ S can ϕ ’ is interpreted as ‘ S can ϕ ψ ly’ – and that postulating such a restriction is not

entirely ad hoc, as the same kind of mechanism arguably explains the context-sensitivity of (15) and (10).²⁹

6 Defending the canon

I now want to briefly revisit some objections that have been raised against the quantificational analysis of ability modals. As we will see, they all involve the transparent sense of ability and are easily answered by the tweaked version of the quantificational analysis that I have outlined.

The first objection might be called the *no-fluke argument*. (It is related to an objection to the conditional analysis that I mentioned in section 2.) On the quantificational analysis, an agent has the ability to ϕ as long as she ϕ s in some accessible world. Often, however, an agent intuitively lacks the ability to ϕ even if her skills and the relevant environmental circumstances are compatible with her ϕ ing. To use the standard example from [Kenny 1976], suppose Betsy is a novice dart player who can barely hit the board. If we look at accessible worlds where Betsy attempts to hit the bullseye, we find the dart landing all over the place. In a few of these worlds (loosely speaking), the dart happens to hit the bullseye. The quantificational analysis therefore seems to imply that (22) is true. But intuitively, it is not.

(22) Betsy can hit the bullseye (on the next throw).

The ability attributed in (22) seems to require a kind of robustness: it shouldn't be a mere fluke that the agent succeeds if she tries.³⁰

In response, we need to distinguish the effective and the transparent reading of (22). On its effective reading, (22) is plausibly true: it is in Betsy's power to do something that amounts to hitting the bullseye. This reading usually isn't salient, but it could be made salient by supplying more context. For example, imagine Betsy has been in an accident and is slowly regaining the ability to move her arm. We discuss how far she could throw a dart in her present state. More precisely, we discuss which points on the wall she could reach if she were to throw a dart from where she is standing. Any point above a height of 3 meters on the wall is ruled out, but points below 2 meters within a certain radius are not. If the bullseye falls in that region, we could truly utter (22).

²⁹ In footnote 5, I mentioned another possible mechanism: scope. I argued that this explanation would not be sufficiently general. But we can now see why scope does play a role – for example, why (6) is intuitively true if 'Eddie's favourite piece' is given wide scope. According to Analysis 2, the wide-scope statement is true (on its transparent reading) iff, informally speaking, there is some x such that (a) x is Eddie's favourite piece, and (b) Maisy has the ability to play x transparently. Is there such an x ? Yes: the Moonlight Sonata.

³⁰ Several alternatives to the quantificational account have been developed around this intuition; see e.g. [Nowell Smith 1960], [Brown 1988], [Greco 2007], [Vihvelin 2013], [Maier 2018b], [Jaster 2020].

(And we would thereby attribute an ability to Betsy – we wouldn’t just express that it could happen that Betsy hits the bullseye.)

In normal contexts, however, (22) is understood transparently, and this reading really does call for a revision to the classical quantificational analysis. On the account I have outlined, (22) requires that there are accessible worlds at which Betsy hits the bullseye transparently. In these worlds, Betsy must have reason to believe that she will hit the bullseye merely on the basis of a relevant decision. By normal accessibility standards, there won’t be any such worlds. So (22) comes out false.

In general, the transparency condition in Analysis 2b ensures that the relevant success events aren’t mere flukes. But note that the no-fluke requirement is epistemic: by Analysis 2b, if an agent has a transparent ability, then *from her own epistemic perspective* it is not an accident if she succeeds. It is not required that the agent succeeds across a whole range of worlds, under varying environmental circumstances. The difference shows up in unusual cases like the following. Suppose Betsy has been informed by a time-traveller that if she will throw the next dart, then she will hit the bullseye. Arguably, that renders (22) true, even on its transparent reading. So a transparent ability to hit the bullseye does not require general skill or success across a range of circumstances.

A second, though related, objection to the quantificational account is what I’ll call the *argument from general abilities*. (22) attributes a “specific” ability, talking about what Betsy can do in her present circumstances. Now we focus on “general” ability statements in which the present circumstances are not held fixed. The objection is that whatever we say about specific abilities, general abilities surely require more than successful performance at a single accessible world.

To illustrate, consider (3): ‘Cyril can recite the first 10 digits of π ’. As mentioned earlier, this is naturally understood as attributing a “general ability”. When we say that someone can recite the first 10 digits of π , we normally mean that they can do so systematically, whenever the need arises – in good weather and in bad weather, in the morning and in the evening.

In response, we again need to distinguish the effective and the transparent reading of the relevant statements. The objection is most plausible for the transparent reading, which is also more salient for sentences like (3). On the account I have outlined, an agent has the transparent ability to recite the first 10 digits of π just in case there is at least one accessible world at which she transparently recites the digits. Suppose there is such a world. In that world, the agent knows that she is reciting the first 10 digits of π merely on the basis of her decision to utter a certain sequence of sounds. That is, she knows that uttering those sounds would amount to reciting the first 10 digits of π . How could she know this? Most likely because she knows what the first 10 digits of π are. So if the transparent reading of (3) is true, then there are accessible worlds at which Cyril

knows the first 10 digits of π .³¹ But whether or not an agent knows the first 10 digits of π is plausibly one of the aspects of the actual world that we tend to hold fixed when we evaluate statements like (3). That is, if Cyril doesn't actually know the first 10 digits of π , then he also doesn't know them in any accessible world.³² The hypothesis that an agent transparently recites the first digits of π in some accessible world therefore implies that she knows what these digits are in *all* accessible worlds. And this, in turn, implies that across the accessible worlds, the agent almost always succeeds to recite the first 10 digits of π when the need arises, irrespective of the weather or the time of the day.

The upshot is that the account I have outlined vindicates the intuition that “general abilities” like the one attributed in (3) imply successful performance across a wide range of trying circumstances, without explicitly putting this into the semantics.

To see why the wide range of circumstances shouldn't be built into the semantics, consider a variant of (7):

(23) Usain Bolt can run 100 meters in under 10 seconds.

It is unlikely that Bolt is *right now* in a position to run 100 meters in under 10 seconds. Perhaps he has only just woken up, has just eaten, is in an airplane, is not wearing his running gear, etc. The sense in which (23) is true (or was true until recently) is therefore a “general” sense. However, unlike (3), (23) does not imply successful performance across a wide range of circumstances. On the contrary, it is enough if there are a few, fairly special conditions under which Bolt can run 100 meters in under 10 seconds. The hypothesis that “general abilities” require success across a variety of circumstances would get (23) wrong. The account I have outlined gets it right: holding fixed Bolt's general physiology and strength, but possibly varying his environment, his state of alertness, etc., we find *some* worlds at which he runs 100 meters in under 10 seconds – including some at which he knows from the outset that he is going to achieve this goal. That is enough to render (7) true, even on its transparent reading.³³

Let me turn to a third objection, due to [Kenny 1976]: the *argument from non-distribution over disjunction*. Imagine a well-shuffled deck of cards, lying face-down on

³¹ This ignores the proviso in Analysis 2b. In fact, according to Analysis 2b, it would suffice for the truth of (3) if (say) Cyril gives credence 0.4 to the digits being 3, 1, 4, 1, 5, 9, 2, 6, 5, 3 and much lower credence to every other possibility. That seems correct to me.

³² As usual, the standards depend on context: ‘Can you recite the first 10 digits of π for this video I'm making?’ – ‘Sure. What are they?’

³³ Other mechanisms might also play a role in generating the apparent robustness requirement in many “general” ability statements. For example, the claim that steel can withstand temperatures of 1000 °C, might be understood to mean that steel can withstand such temperatures *under otherwise normal conditions*, even if the actual utterance context is non-normal. Here a stereotypical ordering source seems to be in play (compare [Kratzer 1981: 60f.]). Or suppose you can only play the piano in a dark and quiet room at temperatures below freezing, and you know that I am looking for a pianist to entertain a group of children. Here it would be misleading to say, without further qualification, that you can play the piano, even if what you say is literally true.

the table. Can you draw a card that's either red or black? Sure: any card will do. But can you specifically draw a red card, in a single try? Arguably not. Nor can you specifically draw a black card. So (24a) seems to be true while (24b) and (24c) are false.

- (24) a. You can draw a red or black card.
 b. You can draw a red card.
 c. You can draw a black card.

The problem is that if 'can' is an existential quantifier over accessible worlds, then 'can(ϕ or ψ)' should entail 'can(ϕ) or can(ψ)': if there are no accessible worlds where you draw a red card, and none where you draw a black card, how could there be accessible worlds where you draw a card that's either red or black?

You can probably anticipate my response. First, (24a)–(24c) are ambiguous between an effective reading and a transparent reading. Since the effective reading is trivial for (24b) and (24c), the transparent reading is more salient.³⁴ And the transparent 'can' really does not distribute over disjunction: there may well be accessible worlds where you ϕ -or- ψ transparently even though there are no accessible worlds where you ϕ transparently nor worlds where you ψ transparently. So (24b) and (24c) can be false while (24a) is true.

Finally, the *argument from hyperintensionality*. Intuitively, (25a) is true while (25b) is false.

- (25) a. Cyril can recite the sequence 3, 1, 4, 1, 5, 9, 2, 6, 5, 3.
 b. Cyril can recite the first 10 digits of π .

The prejacent in (25a) and (25b) are strictly equivalent: there is no possible world at which anyone recites 3, 1, 4, 1, 5, 9, 2, 6, 5, 3 without reciting the first ten digits of π , or the other way round. But if the set of possible worlds at which an agent ϕ s is identical to the set of worlds at which she ψ s, the quantificational analysis implies that the agent has the ability to ϕ just in case she has the ability to ψ . So how can (25a) and (25b) differ in truth-value?

Again, you can probably anticipate my response. On their (non-salient) effective reading, (25a) and (25b) really do stand and fall together. But on their more salient transparent reading, they can come apart. The hyperintensionality comes from the

³⁴ Without further context, arguably neither reading is definitely favoured – especially in the presence of (24a), which is trivial on either reading. This might explain why some authors have resisted Kenny's argument, suggesting that (24b) and (24c) are actually true.

In my own experience, the longer I think about sentences like (24b) and (24c), or indeed (2), (3), or (6), the more I become inclined to say that they are true. For some reason, the effective reading seems to become more salient. If the transparent reading arises through a contextual restriction on the interpretation of the prejacent, this might be explained by the well-known fact that it is generally easier to lift a contextual restriction than to re-impose it.

hyperintensionality of warrant and belief: if Cyril does not know the first digits of π , there can be accessible worlds at which he transparently recites 3, 1, 4, 1, 5, 9, 2, 6, 5, 3 even though there no worlds at which he transparently recites the first 10 digits of π .

7 Conclusion

In the recent literature, a near consensus has emerged that the agentive sense of ‘can’ (that may be paraphrased with ‘able’) defies the classical quantificational account of modals. I have resisted this consensus. On the account I have outlined, ‘can’ always functions as an existential quantifier over possible worlds; different uses of ‘can’ largely correspond to different ways of determining the domain of quantification.³⁵

All else equal, I take it that a unified analysis is preferable. Across a wide range of languages, the same lexical items are used to express agentive modality and other kinds of modality. That can’t be a coincidence. What’s more, the agentive ‘can’ shares many logical and linguistic peculiarities with other uses of ‘can’, such as the way it interacts with ‘if’-clauses or its tendency to trigger “free choice” effects.³⁶ The quantificational account explains these peculiarities. It also unifies the bewildering variety of subflavours within ability modals, as we saw in section 3, and it explains the connection between ability modals and other modals.

The main obstacle for a quantificational treatment of ability modals arises from what I have called their “transparent” reading. To get around this, I have suggested that conversational context may impose restrictions on what counts as a performance of the relevant act among the accessible worlds. On the transparent reading, an agent has the ability to ϕ iff there are accessible worlds at which she ϕ s simply by deciding to ϕ – more precisely, as the result of a decision that warrants believing that she will ϕ , provided ϕ ing is under her volitional control.

Why should we be interested in this peculiar sense of ability? Because it plays a central role in normative contexts, when we consider what an agent in a given choice situation ought to do. Informally, an act is a candidate for being obligatory only if the agent knows what she would have to do in order to perform the act. The transparent concept of ability is tailored to this normative job description.

³⁵ The determination depends on syntactic position: above aspect, ‘can’ tends to have an epistemic reading, below aspect it is circumstantial or agentive. See [Hacquard 2009], [Kratzer 2013].

³⁶ An example of a free choice effect is the inference from ‘I can come tomorrow or on Friday’ to ‘I can come tomorrow’. This is arguably a pragmatic phenomenon akin to a scalar implicature. (Note that ‘I can’t come tomorrow’ does not seem to imply ‘I can’t come tomorrow or on Friday’). As [Nouwen 2018] points out, popular explanations of this effect assume that ‘can’ is an existential quantifier over possible worlds.

References

- G.E.M. Anscombe [1957]: *Intention*. Ithaca, NY: Harvard University Press
- John L. Austin [1961]: “Ifs and Cans”. In *Philosophical Papers*, London: Oxford University Press, 151–180
- Rajesh Bhatt [2006]: *Covert modality in non-finite contexts*. Walter de Gruyter
- Virginia Mary Brennan [1993]: *Root and epistemic modal auxiliary verbs*. PhD Thesis, University of Massachusetts at Amherst
- Mark A. Brown [1988]: “On the logic of ability”. *Journal of Philosophical Logic*, 17(1): 1–26
- Erik Carlson [1999]: “The Oughts and Cans of Objective Consequentialism”. *Utilitas*, 11(01): 91–96
- Robyn Carston [2010]: “Explicit communication and ‘free’ pragmatic enrichment”. In B. Soria and R. Romero (Eds.) *Explicit communication*, Basingstoke: Palgrave Macmillan, 217–285
- Nate Charlow [2013]: “What we know and what to do”. *Synthese*, 190(12): 2291–2323
- B. Jack Copeland [2002]: “The Genesis of Possible Worlds Semantics”. *Journal of Philosophical Logic*, 31(2): 99–137
- Charles B. Cross [1986]: “‘Can’ and the logic of ability”. *Philosophical Studies*, 50: 53–64
- Cian Dorr and John Hawthorne [2013]: “Embedding Epistemic Modals”. *Mind*, 122(488): 867–913
- Michael Fara [2008]: “Masked abilities and compatibilism”. *Mind*, 117(468): 843–865
- Holly S. Goldman [1978]: “Doing the Best One Can”. In A. Goldman and J. Kim (Eds.) *Values and Morals*, Dordrecht: Reidel, 185–214
- John Greco [2007]: “The nature of ability and the purpose of knowledge”. *Philosophical issues*, 17(1): 57–69
- Valentine Hacquard [2009]: “On the interaction of aspect and modal auxiliaries”. *Linguistics and Philosophy*, 32(3): 279
- [2010]: “On the event relativity of modal auxiliaries”. *Natural language semantics*, 18(1): 79–114
- Alison Hall [2008]: “Free enrichment or hidden indexicals?” *Mind & Language*, 23(4): 426–456
- A.M. Honoré [1964]: “Can and Can’t”. *Mind*, 73(292): 463–479
- Terence Horgan [1979]: “‘Could’, possible worlds, and moral responsibility”. *The Southern Journal of Philosophy*, 17(3): 345–358
- Romy Jaster [2020]: *Agents’ Abilities*. Berlin: de Gruyter

- Anthony Kenny [1976]: “Human abilities and dynamic modalities”. In J. Manninen and R. Tuomela (Eds.) *Essays on explanation and understanding*, Springer, 209–232
- Niko Kolodny and John MacFarlane [2010]: “Ifs and oughts”. *The Journal of philosophy*, 107(3): 115–143
- Angelika Kratzer [1977]: “What ‘must’ and ‘can’ must and can mean”. *Linguistics and Philosophy*, 1(3): 337–355
- [1981]: “The notional category of modality”. *Words, Worlds, and Contexts*: 38–74
- [1986]: “Conditionals”. *Chicago Linguistics Society*, 22(2): 1–15
- [1991]: “Modality”. In A. von Stechow and D. Wunderlich (Eds.) *Semantik Handbuch*, Berlin: de Gruyter, 639–650
- [2013]: “Modality for the 21st Century”. In S.R. Anderson, J. Moeschler and F. Raboul (Eds.) *The Language-Cognition Interface*, Librairie Droz., 179–199
- Saul Kripke [1980]: *Naming and Necessity*. Oxford: Blackwell
- David Lewis [1976]: “The Paradoxes of Time Travel”. *American Philosophical Quarterly*, 13: 145–152
- [1979]: “Scorekeeping in a Language Game”. *Journal of Philosophical Logic*, 8: 339–359
- [1986]: *On the Plurality of Worlds*. Malden (Mass.): Blackwell
- John Maier [2015]: “The agentive modalities”. *Philosophy and Phenomenological Research*, 90(1): 113–134
- [2018a]: “Abilities”. In Edward N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, spring 2018 edition
- [2018b]: “Ability, modality, and genericity”. *Philosophical Studies*, 175(2): 411–428
- Matthew Mandelkern, Ginger Schultheis and David Boylan [2017]: “Agentive modals”. *Philosophical Review*, 126(3): 301–343
- James D. McCawley [1978]: “Conversational implicature and the lexicon”. vol 9. New York: Academic Press, 245–259
- Alfred R. Mele [2003]: “Agents’ abilities”. *Noûs*, 37(3): 447–470
- George E. Moore [1912]: *Ethics*. London: Williams and Norgate
- Rick Nouwen [2018]: “Free Choice and Distribution Over Disjunction: the Case of Free Choice Ability”. *Semantics and Pragmatics*, 11. URL <https://doi.org/10.3765/sp.11.4>
- P.H. Nowell Smith [1960]: “Ifs and cans”. *Theoria*, 26(2): 85–101

- Christopher Peacocke [1999]: *Being known*. Oxford: Clarendon Press
- Douglas W. Portmore [2019]: *Opting for the Best: Oughts and Options*. New York: Oxford University Press
- François Recanati [2013]: “Pragmatic Enrichment”. In *Routledge Companion to Philosophy of Language*, Routledge, 89–100
- Kieran Setiya [2008]: “Practical knowledge”. *Ethics*, 118(3): 388–409
- Nicholas Southwood and David Wiens [2016]: ““Actual” does not imply “feasible””. *Philosophical Studies*, 173(11): 3037–3060
- Travis Timmerman and Yishai Cohen [2016]: “Moral Obligations: Actualist, Possibilist, or Hybridist?” *Australasian Journal of Philosophy*, 94(4): 672–686
- Peter van Inwagen [1983]: *An Essay on Free Will*. Oxford: Clarendon Press
- David Velleman [1989]: *Practical Reflection*. Princeton: Princeton University Press
- Barbara Vetter [2013]: “‘Can’ without Possible Worlds: Semantics for Anti-Humeans”. *Philosophers’ Imprint*, 13(16)
- Kadri Vihvelin [2004]: “Free will demystified: A dispositional account”. *Philosophical Topics*, 32(1/2): 427–450
- [2013]: *Causes, laws, and free will: Why determinism doesn’t matter*. Oxford: Oxford University Press
- Ann Whittle [2010]: “Dispositional abilities”. *Philosophers’ Imprint*, 10(12): 1–23