

Ability, modality, and genericity

John Maier¹

Published online: 3 February 2017
© Springer Science+Business Media Dordrecht 2017

Abstract Accounts of ability in the philosophical literature have tended to be modal ones: claims about an agent’s abilities are understood in terms of what she does in certain non-actual scenarios. In contrast, a prominent account of ability ascriptions in the recent semantics literature appeals to genericity: claims about an agent’s abilities are understood in terms of what she generally manages to do. The latter account resolves some long-standing problems for modal accounts, but encounters problems of its own. I propose a hybrid view, on which ability involves both a modal and a generic element.

Keywords Abilities · Genericity · Modal logic

1 Ability and actuality

A philosophical theory of ability should accommodate, at a minimum, the observation that someone may have the ability to perform some action even though she is not now performing that action. For instance, I have the ability to raise my left arm even though I am not now raising my left arm.

A natural way to accommodate this observation is to take ability to be a kind of modality, and to take modality in turn to be understood in terms of quantification over possible worlds. What it is for someone to have the ability to perform some action is for it to be the case that there is some accessible world (or set of worlds) where she performs that action. Call this the *modal view* of ability. On such a view, we have a simple explanation of how someone may have an ability to perform some action that she is not now performing. The modal view requires only that there be

✉ John Maier
john@jmaier.net

¹ Department of Philosophy, Peking University, Beijing 100871, China

some accessible world or other where she performs that action. It does not require that that world be actual.

As stated, the modal view is a very general view, of which more particular views of ability may be regarded as special cases. For instance, several authors have proposed that an agent has the ability to perform some action just in case she would perform that action if she tried to. This is the well-known ‘conditional analysis’—more accurately called the counterfactual analysis—of ability. If we understand counterfactuals themselves in terms of possible worlds, then this is simply a species of the modal view. For instance, when we combine the counterfactual analysis with the semantics for counterfactuals proposed in (Stalnaker 1968), we arrive at the proposal that someone has the ability to perform some action just in case she performs that action at the closest possible world at which she tries to perform that action.

The modal view of ability, then, may take many forms. In any of its forms, it elegantly accommodates the observation that actually performing an action is not *necessary* for having the ability to perform that action. It encounters a difficulty, however, when we consider the question of whether actually performing some action is *sufficient* for having the ability to perform that action. The proper answer to this question is notoriously unclear. Consider:

Tara Tara is a beginner at golf. She misses most of her shots. On this occasion, however, she strikes the ball from the tee and it happens to go in the hole. Tara has, on this occasion, made a ‘hole-in-one’

Does Tara have the ability to make a hole-in-one, or does she not? Here it is tempting to answer equivocally: in one sense she does, and in another sense she does not.

Such equivocation is, at first pass, troubling for the modal view. The trouble may be brought out as follows. Modalities are typically divided into *alethic* and *non-alethic* modalities. For alethic modalities, actuality suffices for possibility. Nomic possibility is a paradigm of an alethic modality. If some state of affairs is actual, then it is nomically possible, for the laws of nature allow at the very least for that which actually occurs. For non-alethic modalities, in contrast, actuality is not sufficient for possibility. If, for instance, we follow many deontic logicians in taking moral permissibility to be a kind of modality, then moral permissibility is a paradigm of a non-alethic modality. From the fact that some state of affairs is actual, it does not follow that it is morally permissible.

A proponent of the modal view of ability then faces the following question: is ability an alethic modality, or is it not? However she answers, she faces another question. If ability is an alethic modality, why are we reluctant to judge without qualification that Tara, who makes the hole-in-one, has the ability to make a hole-in-one? And, if ability is a non-alethic modality, why are we reluctant to judge without qualification that Tara *lacks* this ability? In short, neither the judgment that Tara has the ability to sink a hole-in-one, nor the judgment that she lacks this ability, seems quite right. But if ability is a modality, then exactly one of these judgments should seem right.

I do not suggest that these questions do not have answers. They have been noticed and rigorously addressed by—among others—Cross (1986), Brown (1988), and Horty and Belnap (1995). The formal development of these questions will be pursued below. But I do want to suggest that these questions should at least give us pause about the modal view of ability, and lead us to ask what alternatives to it there might be.

The request for alternatives may seem puzzling. After all, we have acknowledged that someone may have the ability to perform some action even though she is not now performing it. How could this be, it might be asked, unless the modal view of ability is true?

To begin to answer this question, it is helpful to think about habits. The claim that Ann smokes may be true even though Ann is not now smoking. In this sense, just as actually now raising one's arm is not necessary for having the ability to raise one's arm, so is actually now smoking not necessary for being a smoker. Yet most do not conclude that claims about habits are modal claims. A standard view is rather that claims about habits rather involve a distinct mechanism—much studied by linguists, and increasingly by philosophers as well—of *genericity*.¹

As it happens, a generic view of ability claims has been prominently defended in the linguistics literature by Bhatt (1999). Bhatt's view is directed at explaining a puzzling linguistic phenomenon—the so-called 'actuality entailment'—but it turns out to have substantive philosophical consequences as well. In particular, it gives us a concrete model of a genuinely amodal view of ability. In what follows I give a more extensive exposition of Bhatt's view and its philosophical consequences.

We therefore have, at least, a choice between a modal theory of ability, on which the crucial mechanism is quantification over possible worlds, and an amodal theory of ability, on which the crucial mechanism is genericity. The virtues and vices of these theories will be the topics of Sects. 2 and 3, respectively. We may also give a hybrid theory of ability, on which ability involves both a modal and a generic element. This is the kind of view that I will, in Sect. 4, defend.

2 The modal view of ability

Consider first the modal view of ability. Certain difficulties for the modal view were first noticed by Kenny (1975). To present Kenny's concerns, it will be helpful to make a distinction. On the modal view of ability, I have said, an agent has the ability to perform some action just in case there is some accessible world (or set of worlds) where she performs that action. Let us distinguish between a view that appeals to a

¹ It may be that modality and genericity are not entirely distinct, for many linguists propose to give a modal account of genericity; see Carlson and Pelletier (1995: 49–57) for citations and discussion. If such views are correct, then a generic view of ability is, at a more fundamental level, itself a kind of modal view. For present purposes I remain neutral on the proper account of genericity; when I say that a generic view is 'amodal' this should be understood only in the uncontroversial sense that there is, at the descriptive level, a distinction to be drawn between modal expressions and generic expressions (even if the latter ultimately admit of some kind of modal analysis).

single world, which I will call a *simple* modal view, and a view that appeals to a set of worlds, which I will call a *complex* modal view. Kenny's objections are brought against the former kind of view. I begin by considering the objections that Kenny brings against such a view; I then consider whether these objections might be answered by turning to a complex modal view, as is proposed in Brown (1988), as well as Cross (1986); finally, I consider whether they might be addressed within the framework of 'stit' logic, as is proposed in Horty and Belnap (1995).

2.1 A simple modal view

The general form of Kenny's arguments is as follows. In light of work in modal logic (notably, Kripke 1963), we should expect that, if a simple modal view of ability is true, ability claims should have certain logical properties. But, Kenny observes, they do not in fact have these properties. So, he concludes, we should hold that a simple modal view of ability is false.

One of the logical properties considered by Kenny is that which corresponds to the **T** axiom in standard modal logics:

$$\mathbf{T} \quad p \rightarrow \Diamond p$$

If we accept a simple modal view of ability, this comes to the claim that if someone performs some action then she has the ability to perform that action. Kenny claims, this is not plausible: 'A hopeless darts player may, once in a lifetime, hit the bull, but be unable to repeat the performance because he does not have the ability to hit the bull' (Kenny 1975: 136).

This problem is essentially the same as the one at issue in the case of Tara. And here it is natural to make one of the responses suggested there: we should simply deny that ability is an alethic modality. This is just to say, in the framework of Kripke (1963), that it fails to satisfy **T**. And indeed, within that framework, **T** is an optional axiom. There are sound and complete modal logics that satisfy it, and sound and complete modal logics that do not. In the language of modal logic, then, our response to this problem should simply be that the logic of ability is of the latter kind, and that it does not obey the **T** axiom.

This would not, as already noted, address our equivocation about the case of Tara: if **T** is simply invalid for ability, then why are we inclined to say that, in some sense, Tara does have the ability to sink-a-hole in one? This is a question to which we will return. However we resolve this question, Kenny notes that there remains a distinct, and more fundamental, logical challenge to the modal view of ability. The following is an axiom of the minimal modal logic of (Kripke 1963):

$$\mathbf{K} \quad \Diamond(p \vee q) \rightarrow \Diamond p \vee \Diamond q$$

Unlike **T**, **K** is not an optional axiom within the framework of the Kripke semantics; it is an axiom of any 'normal' modal logic. Yet, as Kenny notes, ability appears not to obey this axiom (Kenny 1975: 137). Consider:

Kara Kara is given a randomly shuffled deck of cards, containing (as a normal deck does) an equal number of red and black cards. She is asked to draw one card from the deck

Here, Kara has the ability to pick a red card or a black card. But, Kenny claims, she does not have the ability to pick a red card, and she does not have the ability to pick a black card. If this claim is correct, then ability does not obey a principle that is axiomatic in a normal modal logic.²

Kenny's response to this point is simple, and drastic: 'if we regard possible worlds semantics as making explicit what is involved in being a possibility, we must say that ability is not any kind of possibility' (Kenny 1975: 140).

2.2 Complex modal views

Kenny's argument, however, presupposes that a modal view must be simple in the sense given above. Kenny does not consider what we have called a *complex* modal view, on which accessibility relations hold, not between an individual world and an individual world, but between an individual world (the world where the ability claim is evaluated) and a set of worlds. Precisely this point is made by Mark Brown, who writes that Kenny's 'pessimism is a bit hasty. The general idea of a possible-worlds semantics is a very flexible one: there are other ways in which we can give a semantics based on possible worlds' (Brown 1988: 3). In particular, Brown proceeds to show, we can give a complex modal view of ability.

Briefly stated, Brown's idea is the following. Say that an agent is able to perform some action just in case there is a relevant *set* of accessible worlds where she performs that action. It can be shown that, as on the modal logic of Kripke (1963), the **T** axiom need not be valid on such a logic, and so we need not accept that Tara has the ability to sink the putt. The advantage of Brown's view comes when we consider the **K** axiom, which need not be valid on such a view either. Say that Kara has the ability to pick a red card or a black card. Then, on Brown's view, at every world in the relevant set she picks a red card or she picks a black card. But it is not the case that at every one of these worlds she picks a red card, nor that at every one of these worlds she picks a black card. Rather, she picks a red card at some of the worlds in the set, and a black card at other worlds in the set. So it is not the case that, for either one of these kinds of cards, Kara picks it at every world in the relevant set. Thus we cannot derive, from the claim that Kara has the ability to pick a red card or a black card, the conclusion that she has the ability to pick a red card, nor that she has ability to pick a black card.

But Brown's view fails to solve one of the problems raised above, and it gives rise to still another. These problems in turn point away from a modal view of ability, be it simple or complex, and towards another kind of view altogether.

² 'Normal' here is being used simply to denote those modal logics that obey the minimal axioms of Kripke (1963). There are 'non-normal' modal logics on which **K** is not valid; we will consider how to implement such a logic to give an account of ability in the next subsection.

The problem already raised above is our equivocation in the case of Tara. I have argued that, whether or not an account claims that Tara has the ability to sink a hole-in-one, it owes us an explanation of our reluctance to unqualifiedly endorse or reject that claim. Brown's view, on which she simply does *not* have the ability to sink a hole-in-one, owes us an explanation of why we are still partly inclined to judge that she does. J.L. Austin writes, of such a case: 'it follows merely from the premise that he does it, that he has the ability to do it, according to ordinary English' (Austin 1956: 218). Brown's proposal gives us no account of this aspect of ability.³

The second problem concerns a different kind of case, one which appears to show that, whatever we say about the case of Tara, the truth-conditions suggested by Brown are far too demanding. Brown writes that his proposal is intended to capture the sense in which an agent who has an ability to A 'can (reliably) bring about circumstances in which A is true' (Brown 1988: 6). But someone may be reliable without it being the case that she brings about these circumstances in *every* world, as Brown's proposal demands. Consider:

Gina Gina is an excellent golfer. When she is confronted with a short putt, she almost always – but not quite always – sinks it. She is now confronting a short putt

It is plausible that Gina has the ability to sink this putt. But it is not the case that she sinks the putt in *every* world in the relevant set of accessible worlds. So Brown's proposal yields the implausible conclusion that Gina does not have the ability to sink the putt. And, if Gina does not have the ability to sink the putt, then many of us lack abilities that we take ourselves to have. We often take ourselves to be able to perform actions that we cannot quite *guarantee*. Brown's view would render many of our ordinary claims false.

At this point, it is tempting to return to the idea that ability claims involve existential quantification after all. Perhaps having an ability to perform an action is a matter of there being a set of worlds such that one performs that action in *some* world in that set. This is the idea developed in Cross (1986). Roughly, Cross proposes that an agent has an ability just in case in the set of worlds that are fortuitous for the exercise of the ability in question (those that satisfy what Cross calls 'appropriate test conditions'), the agent performs the action in one of those worlds. So, since Gina indeed satisfies this condition, Gina has the ability to sink a putt.

There is not space here for a full treatment of Cross's elegant proposal, but the basic concern about it will be clear. While Cross's view is not a simple modal view (see Cross 1986: 60–61), it shares with a simple modal view the claim that Tara is simply able to make the putt. More generally, it seems to predict that an agent who would successfully perform an action once in appropriate test conditions—even if she never *actually* succeeds, and even if she would fail almost every other time—

³ Brown does introduce an operator that existentially quantifies within sets of worlds, which he describes as a *might* operator (Brown 1988: 7). So he can acknowledge that, though Tara is not able to sink a hole-in-one, she might sink a hole-in-one (since, after all, she does). A formally similar operator is developed, and proposed as an account of ability, in Cross (1986), which I return to below.

has the ability to perform that action. But that renders ability claims excessively undemanding.

We therefore face a dilemma. When we develop a complex modal view in terms of universal quantification, as does Brown, ability claims turn out to be too demanding. We develop a complex modal view in terms of existential quantification, as does Cross, ability claims turn out to be not demanding enough. What we seem to need is a kind of quantification which neither existential nor universal quantification is quite apt to capture: what is required is more than one instance, but less than every instance. It is precisely such phenomena that have led linguists to appeal to *genericity*. To anticipate, I will propose that an appeal to the generic element in ability is precisely the way out of the dilemma.

2.3 The stit view

Before coming to that, it bears considering a different way of modeling ability within the framework of modal logic, one developed within the ‘stit’ theory originally proposed in Belnap and Perloff (1988). Whereas Brown and Cross give a more complicated account of the accessibility relation, this approach retains the thought that ability claims involve a simple possibility operator. The distinctive logical behavior of ability claims in cases like Tara and Kara is to be explained instead in terms of the logical behavior of what occurs under the scope of this operator. While the stit view of ability is strictly a version of a simple modal view, it ends up being quite close in its predictions to the kind of view developed by Brown.⁴ And, I will argue, it is subject to the same objection.

Stit theory is distinguished by its introduction of an operator, *stit* (abbreviating ‘sees to it that’), which relates agents and propositions. Thus the formula:

S stit A.

Abbreviates: ‘S sees to it that A’. This theory is constructed in terms of a framework of branching histories: an agent sees to it that A, at some moment *m*, just in case she does something such that A is the case at *every* history that includes *m*.⁵ For instance, if I make coffee for the meeting at some moment, I thereby make it the case that every future history is—whatever else may be true at it—one at which coffee is made for the meeting.

One advantage of this account is that it allows us to give an elegant account of ability claims. How shall we represent the claim that I am *able* to make coffee for the meeting? As follows:

⁴ See Horty and Belnap (1995: 611–615) for independent demonstration of the similarity between Brown’s proposal and the stit proposal.

⁵ The literature on ‘stit’ has distinguished between several ‘stit’ operators, and the one described here is what is called the ‘deliberative stit’ or ‘dstit’ operator. See Horty and Belnap (1995: 587–595) for an explanation of the difference between this operator and the operator originally introduced in Belnap and Perloff (1988). The treatment of Kenny’s cases given here follows the presentation in Horty and Belnap (1995: 610–611); see also Horty (2001: 20–21).

$\Diamond S$ stit A

Here ‘S’ is the agent in question and A the proposition that coffee is made for the meeting. ‘ \Diamond ’ is a simple alethic modal operator, satisfying both the **T** and the **K** axioms. We have already seen Kenny’s objections to a simple modal view. But the distinctive aspects of the stit theory allow this approach to respond to these objections.

Consider Tara. On the present view, the actual world is accessible to itself, so there is a possible world at which Tara sinks the hole-in-one. Thus, letting ‘A’ be the proposition that Tara sinks a hole in one, ‘ $\Diamond A$ ’ is true. But it is not the case that, at the actual world, Tara sees to it that she sinks a hole-in-one. For, being a poor golfer, she does not do something (at the moment of striking the ball) such that she sinks the hole-in-one at *every* history; she at best does something such that she sinks the hole-in-one at *some* history, namely the one that turns out to be actual. So even though, on the present view, the actual truth of a proposition entails its possible truth, we cannot derive ‘ $\Diamond S$ stit A’ from Tara’s sinking the hole-in-one, for ‘S stit A’ is not true in the first place. Thus it is not the case, on the proposed account of ability, that Tara has the ability to sink a hole-in-one.

The case with Kara is much the same. It is possible that she sees to it that she picks a red card or picking a black card. But it is not possible that she sees to it that she picks a red card, nor that she sees to it that she picks a black card. The most that Kara can do is to pick a card from a deck containing red and black cards—she thereby sees to it that she picks a red card or a black card, but she does not thereby see to it that she picks any color of card in particular. So Kara has the ability to pick a red card or a black card, but she does not have the ability to pick a red card, and does not have the ability to pick a black card. In short, what is true of Kara is:

 $\Diamond S$ stit (A \vee B)

But from this it does not follow that:

 $\Diamond S$ stit(A) \vee $\Diamond S$ stit(B)

Note that this is so even though, on the present view, the **K** axiom is valid for possibility: the possibility operator distributes over disjunction just as it does on any simple modal view. The crucial point is that the initial ability claim about Kara does not involve a possibility operator scoping over a disjunction; rather, the disjunction occurs within the scope of the stit operator itself.

The stit approach has a further virtue, one which gives it at least one advantage over the complex modal view defended by Brown. I objected to Brown that he could not accommodate our equivocation in the case of Tara. In contrast, the stit theory has a nice explanation of this fact. As we have seen, while it is not true that it is possible that Tara sees to it that she sinks the hole-in-one, it is true that it is possible that Tara sinks the hole-in-one (since she actually sinks it). Our equivocation arises then from an equivocation between two possibility claims, one of which ($\Diamond A$) is true, but the other of which ($\Diamond S$ stit A) is false.

Yet the stit theory of ability, at least in the form given, should be rejected. For it does not deliver a plausible answer in the case of Gina. Given our description of the case, it is not possible such that Gina does something (at the moment of striking the

ball) such that she sinks the putt in *every* history. At best, it is possible that she does something such that she sinks the putt in *almost every* history. But then it is not true of Gina, on the stit theory, that she has the ability to sink the putt. And, as I have already objected to Brown, this counterintuitive result threatens to ramify widely. We often take ourselves to be able to perform actions that we cannot quite *guarantee*.⁶ The stit view, like Brown's view, would render many of our ordinary claims false.⁶

3 The generic view of ability

The modal view of ability has been adopted in semantics as well as in philosophy. If we take 'able to' to be a modal auxiliary verb, and if we adopt the semantics for such verbs proposed in Kratzer (1981), then we give a modal semantics for a sentence such as:

(1) Sara is able to swim to the beach

On Kratzer's view, (1) is true just in case there is a possible world w such that: (i) Sara swims to the beach at w , (ii) w is compatible with some contextually-specified set of facts F , and (iii) w is sufficiently close to the actual world relative to some contextually-specified 'ordering source'.

Bhatt (1999) notes a purely semantic problem for this approach. Consider:

(2) Sara was able to swim yesterday

Intuitively, there are two readings of (2). On one reading, (2) simply reports that Sara yesterday had the ability to swim. But, on another reading—arguably the more natural reading in non-philosophical contexts—(2) reports that Sara *did* in fact swim yesterday. On this reading, (2) is true only if Sara actually swam yesterday. In Bhatt's term, (2) on this reading has an *actuality entailment*.

The actuality entailment is not overtly marked in (2), but this is an accident of English. Other languages (those that have distinct imperfective and perfective morphology) do mark this entailment. Consider:

⁶ Note that, in a way, the problem of Gina is even more general than this. For consider the case where Gina *actually does* make a short putt. Still, in such a case, Gina is not able to sink a putt, for she does not see to it that she sinks a putt—for there are a few non-actual histories on which she does *not* sink a putt. Therefore it is not the case that Gina sees to it that she sinks a putt, and more generally it is never the case that, in similar circumstances, Gina sees to it that she sinks the putt. Therefore it is not the case \Diamond stit A. In this sense, this modified case of Gina is not distinguished, within the stit theory, from the case of Tara.

One might suggest that what the case of Gina shows is not a special problem about ability but a problem about the generic aspect of action itself, a phenomenon emphasized in Thomason (ms.). This, I think, would be a mistake. There are treatments of action claims (such as Davidson 1967) that can adequately handle the fact that agents sometimes perform actions without guaranteeing the outcomes of their actions. If the stit theory cannot accommodate that fact, then that is potentially an independent objection to the stit theory. But whatever our treatment of action claims, we need an account of *ability* claims that appropriately handles the case of Gina. That is what, I am arguing, neither the stit view nor Brown's view successfully delivers.

- (3) Sara a pu nager hier
Sara was able to swim yesterday
- (4) Sara pouvait nager hier
Sara was able to swim yesterday

(4) here demands only that Sara yesterday had the ability to swim. (3), in contrast, bears an actuality entailment: it is true only if Sara actually swam yesterday.

(3) is a problem for the Kratzerian view of ability ascriptions described above. On the standard application of that view, (3) should demand only that there be some possible world—which must satisfy certain specifications but which need not be actual—where Sara swam yesterday. But (3) demands more than this. As just noted, it is true only if Sara actually swam yesterday. We therefore face a problem for the modal view distinct from the more philosophical considerations brought against it in the previous section.

Bhatt proposes, in light of this problem, that we should distinguish two senses of ‘is able to’. The basic sense is that involved in (3) and in the latter reading of (2). In this sense, ‘able to’ simply behaves like an implicative verb, roughly synonymous with ‘manage to’. The secondary sense of ‘is able to’—the one involved in (4) and the former reading of (2)—is given by a generic operator (*Gen*) over such a verb. Informally, the idea is that ‘is able to’ in English is ambiguous: on one reading, it has the meaning of ‘manage to’; on its other reading, it has the meaning of ‘generally manages to’.

Bhatt’s proposal straightforwardly derives the actuality entailment. Someone manages to perform some action only if she actually performs it. So if one of the meanings of ‘is able to’ is roughly like that of ‘manage to’, it is unsurprising that it requires that the agent actually perform the action in question. And since it is the only available reading, on Bhatt’s proposal, of (3), we thereby explain why (3) bears an actuality entailment.

Bhatt’s proposal is a semantic one, a theory of ability ascriptions. But if we allow ourselves to move from the formal to the material mode, it yields a theory of ability as well. On its simplest development, the view can be stated as follows. There are in fact two kinds of ability—ability₁ and ability₂—corresponding to the two senses of ‘is able to’. Someone has the ability₁ to perform an action just she manages to perform that action; someone has the ability₂ to perform an action just in case she generally manages to perform that action, where ‘generally’ is stipulated to have the semantics of the generic operator (*Gen*). Call this the *generic view* of ability. The generic view, as it turns out, deals quite well with the problems brought against the modal view in the previous section.

Consider first Tara. Above I noted that our response to the question of whether Tara has the ability to sink a hole-in-one was equivocal, and that the modal view, simple or complex, could not accommodate our equivocation. The generic view, in contrast, traces our equivocation to an ambiguity in the question itself. There are really two questions here: does Tara have the ability₁ to sink a hole-in-one, and does she have the ability₂ to sink a hole-in-one? Tara does, on the generic view, have the ability₁ to sink a hole-in-one: after all, she does manage to sink a hole-in-one. She

does not, on the other hand, have the ability₂ to sink a hole-in-one: she does not generally manage to sink a hole-in-one. This is why our answer to the question of whether she has the ability to sink a hole-in-one simpliciter is precisely: in one sense she does, and in another she does not.

Consider then Kara. Here, since we are considering a case where Kara does not actually pick a card, ability₁ is not relevant: Kara has the ability to pick a certain kind of card just in case she has the ability₂ to pick that kind of card. And the generic view delivers the intuitively correct verdict in this case. Kara does have the ability to pick a red or a black card, for she generally manages to pick a red card or a black card. She does not, however, have the ability to pick a red card, nor the ability to pick a black card. For she does not generally manage to pick a red card, nor generally manage to pick a black card: sometimes she picks one, and other times she picks the other. This example points to a fundamental difference between the logic of possibility and that of genericity. In standard modal logic, as we have said, a disjunctive proposition is possible only if one of its disjuncts is possible. In contrast, this principle is invalid for genericity. For instance, newborn babies are generally male or female, but they are not generally male, and they are not generally female. The fact that one can have the ability to perform a disjunction without having the ability to perform either of its disjuncts is, on the generic view, simply an instance of this aspect of genericity.

Bhatt's amodal view, then, can handle the case of Kara; so too, as noted in the previous section, can Brown's view and the stit view. But, as we noted above, these views are subject to two further objections. First, Brown's view cannot explain our equivocation in the case of Tara. In contrast, the generic view, as we have just seen, explains this equivocation (as did the stit view). Second, and crucially, both Brown's view and the stit view make ordinary ability claims excessively demanding. It is not always the case that someone with the ability to perform an action performs it in *all* of the relevant worlds, as Brown's view demands; nor is it the case that it is possible that she do something that will see to it that she perform that action in *every* subsequent timelines, as the stit view demands. Our abilities, even when reliable, are not fail-safe. Gina may have the ability to make a putt so long as she makes it in *most* of the relevant worlds. Neither Brown's view nor the stit view, it was objected, could not accommodate such considerations.

The generic view, in contrast, easily accommodates them. Gina does not always or invariably make a short putt, but she does generally manage to make a short putt. It is in virtue of this fact, on the generic view, that she has the ability to make a short putt. For what it is to have an ability to perform an action on the generic view—in the relevant sense of 'ability'—is simply to generally manage to perform it.⁷

⁷ Until we have an adequate semantics of the (*Gen*) operator—something I am not sure we in fact have—this solution remains, in part, a promissory note. It may even be, as noted above, that the (*Gen*) operator itself ultimately admits of a modal analysis. But this is no objection per se to Bhatt's account, or to the modified generic account that I will eventually defend. These accounts shed light on the problems posed by ability by either solving these problems or, as in this case, by reducing them to more general outstanding problems, in this case to the semantics of (*Gen*) itself.

So, in addition to accommodating the purely semantic considerations that initially motivate it, the generic view solves the problems brought against the modal view of ability in the previous section. Yet the generic view faces an obvious, and seemingly decisive, objection.

The objection concerns the many actions that each of us is able to perform but never performs, and never will perform. I, for example, have not and (let us stipulate) never will read *War and Peace*. I therefore have not managed and will not manage to read *War and Peace*. Nor do I generally manage to read *War and Peace*. Yet (it seems) I am nonetheless able to read *War and Peace*. The generic view seems to have difficulty explaining the truth of this simple ability ascription.

Bhatt notes this problem, but his response to it is brief and unconvincing. Bhatt notes that ‘not all generic sentences require verifying instances’ (Bhatt 1999: 186), giving some familiar examples from the literature on generics, such as:

- (5) This machine crushes up oranges and removes the seeds
- (6) Sally handles the mail from Antarctica

(5) and (6) plausibly involve genericity, and plausibly are true even if the machine is never switched on and no mail ever arrives from Antarctica, respectively. Such cases need to be accommodated by a complete theory of generics, but they are comparatively marginal. But in ability ascriptions, the absence of verifying instances appears to be the normal, or at least a predominant, case. There are an indefinite number of true ability ascriptions without any verifying instances whatsoever. Equally, there are an indefinite number of false ability ascriptions without any ‘falsifying’ instances whatsoever. For instance, it is false that I have the ability to read *War and Peace* in the original Russian. The generic view owes us a mechanism to explain why it is true that I have the ability to read *War and Peace*, while it is false that I have the ability to read *War and Peace* in the original Russian, though I have never (and will never) do or try to do either. And it needs to do this without invoking that which is to be analyzed, namely my abilities themselves. And Bhatt gives us no indication of how this might be done.⁸

4 The hybrid view of ability

The generic view of ability elegantly handles some central problems for the modal view, but faces a problem that—absent some further mechanism—appears to be decisive. But we should not discard Bhatt’s insights, and I will now propose a theory of ability that builds upon them. This proposal will adopt two of Bhatt’s main ideas. First, ability claims have two senses. Second, the latter of these senses is derived

⁸ A generic account of ability ascriptions which promises to overcome this problem (and others) has recently been proposed in Mandelkern et al. (ms.). There is not space here to discuss this rich proposal, but it bears emphasizing that it is developed in terms of the classical idea that an agent’s abilities are to be understood in terms of conditionals relating her volitions to her actions. The account that I will propose in the next section, in contrast, makes no such assumption, instead treating a certain species of agentive modality as non-conditional and indeed as analytically primitive.

from a generic operation on the first of them. The difference will lie in its account of that over which the generic operator operates. Whereas for Bhatt this is an implicative verb, on the alternative proposal this is a relation that I will call, following (Maier 2015), an *option*. Options are modal, at least in the sense that they have a modal entailment: someone has the option to perform an action only if there is a possible world where she performs that action. In this sense, the proposal to be developed combines the generic aspects of Bhatt's view with the modal aspects of the standard philosophical view, and inherits the virtues of each.

4.1 Foundations: options

Abilities, then, are to be understood in terms a generic operator over options. So what are options?

To begin, consider an agent, attempting to decide what to do *here and now*. She is interested in, for example, how to get to work. Such an agent confronts a variety of actions, only some of which are open to her. For instance, she may take a bus at a certain cost, or a train at a higher cost; these actions are open to her. But, if there are no helicopters available, she may not take a helicopter to work, at whatever cost; taking a helicopter to work is not open to her, at any cost. Such an agent's *options*—and, mutatis mutandis, any agent's options—are simply all and only those actions that are in this sense open to her.⁹

An agent's options are not to be identified with her abilities. An agent may, for instance, have the ability to take a bicycle to work, but this will not be among her options if no bicycles are available. An agent's abilities do in some sense constrain her options: an agent who lacks the ability to ride a unicycle typically does not have the option of riding a unicycle to work, whether or not there are any unicycles available. So having an option and having an ability are two distinct—though related—matters.

The notion of an option, we may say, is *particular*: it depends on facts about the agent and her immediate environment in the situation as it actually is, and not as it might be on some other occasion. I will also take it to be *objective*: an agent's options—as I will be using the term 'options'—do not depend on what actions she believes (reasonably or not) to be open to her. For instance, if an agent confronts a glass, she has the option of drinking from that glass; if the glass is full of gin, then she has the option of drinking gin, even if she believes (however reasonably) the glass to be full of gasoline.

In short, a description of an agent's options—in this stipulative, but I hope not unfamiliar, sense of 'option'—is what we might call a *realistic* description of her choice situation.¹⁰ It describes the actions that are actually available to her, whether

⁹ The 'open to' here locution is not offered as an *analysis* of what is meant by options, but only as an elucidation in somewhat more familiar terms; for a defense of the idea that options are analytically primitive, see Maier (2015: 119–120).

¹⁰ While the appeal to options is partly inspired by their centrality to decision theory, the term 'option' is at this point being used in a way that departs from the way it is often used in descriptions of decision theories. Sometimes it is assumed that an agent's options are transparent to her, in the sense that if an

or not they would be the same on some other occasion, and whether or not they are the same as the options that she takes herself to have.

4.2 Principles for options

Having fixed on this notion of having an option, we can bring it to bear on the case of Tara. Does Tara have the option of sinking a hole-in-one, in the case described? She does—she actually did it, so it must have been among her options. A description of her choice situation that omitted sinking a hole-in-one would have omitted an option that, as it turns out, she in fact had. Of course, there are other things we may say about this option—she did not know that she had this option, she does not reliably have this option—that may have made disregarding it appropriate for a variety of purposes. But if our question is whether she had the option to perform this action simpliciter, then the answer is that she did.

I propose, then, that options satisfy the following principle:

Actuality If S does A, then A is an option for S

It is in virtue of the fact that options satisfy **Actuality**, I will eventually argue, that we judge that in one sense—but only in one sense—Tara has the ability to sink a hole in one.

Though actually performing an action is sufficient for having that action as an option, it is not necessary. An agent may have the option of performing actions that she never in fact performs. I propose that this aspect of options is in fact a modal one. Options, I propose, obey also the following principle.

Possibility If A is an option for S, then there is a possible world where S does A¹¹

Footnote 10 continued

action is an option for an agent then she knows (or is in a position to know) that it is an option for her. Options as they are being used here are not subject to any such epistemic requirement. In some respects, what I here call 'options' are more similar to what are typically called 'outcomes' in decision theory. For instance, the decision theorist's 'outcomes' satisfy the principle that of **Actuality** that I will presently introduce. In other respects they are different. A full treatment of the relationship between the framework adopted here and the ontology of standard decision theory is beyond the scope of the present discussion.

Note also that, despite my earlier rejection of the stit theory of ability, there are important similarities between the notion of an option that figures here and the notion of an 'action' (or 'choice cell') as it figures in the stit theory. For an important recent treatment of this aspect of stit theory, as well as a modification of it intended to allow for an epistemic requirement on actions, see Horty and Pacuit (ms).

¹¹ In recent work (Spencer ms.), Jack Spencer devises a number of ingenious cases that purport to show that agents have the ability to perform actions such that it is metaphysically impossible that they perform them (see also Fara 2010). If Spencer's cases are sound, they pose a further problem for modal analyses of ability. For present purposes, it suffices to note that these cases do not show that agents ever have the option of performing such actions; as such, they are compatible with Possibility.

There is an affinity between this principle and the simple modal analysis of ability. Indeed, I believe **Possibility** captures the core of truth in that view. But there are two important differences. First, as noted, this is not an analysis of ability, but rather concerns the distinct, though related, notion of an option. In the next section, I will explain how we may give an analysis of ability in terms of options. Second, this is not a modal *analysis* of options. From claims about agents' options, we can derive claims about possibility. But the converse does not invariably hold.

This point is crucial when we consider the case of Kara. Kara, it is clear, has the option of picking a red card or a black card. She need not, however, have the option of picking a red card, nor the option of picking a black card. The present view can accommodate these facts. Given what we have said, conjoined with **Possibility**, there is a possible world where Kara picks a red card or a black card. Since she picks a red card or a black card at that world, there is therefore a possible world where she picks a red card, or a possible world where she picks a black card. But from this it does *not* follow that she has one or the other action as an option. That would be true only if we were to give a modal analysis of options. But the case of Kara shows precisely, I think, that we should not do this. Options are modal in the sense of entailing claims about possible worlds, not in the sense of being fully analyzable in terms of them. I return to the philosophical consequences of this point in the conclusion.

4.3 The hybrid view stated

To have an option is not to have an ability. Someone may have the ability to perform some action without now having that action as an option. For instance, a skilled golfer may lack both club and course and so, while retaining her ability to play golf, nonetheless not have playing golf as an option. Though options and abilities are distinct, they are not unrelated. In particular, as I will now argue, we may give an account of ability in terms of options.

What is it, then, to have an ability? I propose: *S has the ability to A just in case A is generally an option for S*. Our definition here contains two terms of art: 'generally' and 'option'. The former is being used stipulatively to refer to the operation referred to in the semantics literature by the generic operator (*Gen*). The latter is being used in precisely the sense in which it was explained in the previous two subsections.

Consider, once more, Tara and Kara. Does Tara have the ability to sink a hole-in-one? No: while she sank the hole-in-one on this occasion, and so has the option of sinking a hole-in one, she does not generally have the option of sinking a hole-in one. Does Kara have the ability to pick a red card, or the ability to pick a black card? No: while she generally has the option of picking a red card or a black card, and so has the ability to pick a red card or a black card, she does not generally have the option of picking a red card, and does not generally have the option of picking a black card. Sometimes she has the one option, and sometimes the other, but neither is such that she generally has it as an option.

Consider then Gina. Does Gina have the ability to sink a putt? Yes: she generally has the option of sinking a putt. The fact that she occasionally lacks the option of

sinking a putt, even in otherwise fortuitous circumstances, does not falsify this claim, just as the fact that some cats have three legs does not falsify the claim that cats (generally) have four legs.

It is its treatment of these kind of exception-admitting ability claims—which, as I have emphasized, are endemic in our thinking about agency—that constitutes the main advantage of the present account of ability. Unlike the simple modal view or Cross's view, the present view recognizes that Tara does not have the ability to sink a hole-in-one. Unlike Brown's view or the stit view, however, it recognizes this fact without thereby having to deny that Gina has the ability to sink a putt. For the kind of universal quantification that gave rise to that problem in those accounts is nowhere present in the account proposed here. Finally, unlike Bhatt's view—which can accommodate the fact that Tara lacks the ability to sink a hole-in-one while Gina has the ability to sink a putt—the present account can recognize the fact that each of us has countless abilities to perform actions that we will never perform, even highly unlikely and unusual actions. For, unlike Bhatt's, it does not try to accommodate ability entirely in terms of genericity and an amodal notion of managing to perform an action. Rather, it appeals to the admittedly modal notion of an option. I have the general ability to read *War and Peace* in virtue of the fact that I generally have the option of doing so; this is so even though I never, in fact, manage to read *War and Peace*. In contrast, I lack the general ability to read *War and Peace* in the original Russian in the virtue of the fact that I do not generally have the option of doing so.

The present proposal is a metaphysical one, an account of what ability is. But it has a natural semantic correlate, an account of the contribution of 'is able' to the sentences in which it occurs. The account is isomorphic to Bhatt's. On this view, a sentence of the form 'S is able to A' has two readings.¹² On the one reading it is true just in case *Gen*(S has the option of Aing). On the other reading, which occurs in contexts where the appeal to genericity is for cognitive or linguistic reasons suppressed, it is true just in case S has the option of Aing.¹³ If has the ability and the option of performing an action, then the claim that she is able to perform that action will be clearly true. If she has neither, then such a claim will be clearly false. If she has one but not the other, then it will be true in some contexts and false in others; absent the clear specification of a context, its truth value may simply be unclear. And that is precisely what happens in the case of Tara.

¹² In short, the two proposed senses for 'able' may be distinguished as follows, letting S be an agent, A an action, and OPT the relation of having an action as an option:

- (i) $[[S \text{ is able to } A]] = 1 \text{ iff } \text{Gen}(S \text{ OPT}(A))$
- (ii) $[[S \text{ is able to } A]] = 1 \text{ iff } S \text{ OPT}(A)$

¹³ These contexts include those which originally motivated Bhatt's proposal. When 'able' is combined with past tense and perfective aspect, the generic reading appeals not to be available. This gives rise to the problem of the 'actuality entailment'. Why should a sentence such as our earlier (3) be acceptable only if the agent actually performed the act in question? Since, on the present proposal, agents often have the option of performing actions they do not perform, the present semantics, unlike Bhatt's, gives no explanation of the actuality entailment. So the development of a treatment of the actuality entailment within the present semantics has to be left as an unsolved problem. For the difficulties faced here by any proposal, including Bhatt's own, see the discussion in Portner (2009: 203–213).

5 Conclusion: the modality in ability

Most philosophers have proposed to give a modal analysis of ability. As indicated at the outset, the well-known ‘counterfactual analysis’ is merely a specific implementation of this more general idea, which has been pervasive in logic and linguistics as well as in philosophy proper. Following Bhatt, I have proposed that this idea is basically mistaken. The operator central to an account of ability is a generic operator, rather than a quantifier (existential or universal) over possible worlds.

Unlike Bhatt, I have proposed that there is a modal element in ability: what the generic operator operates over is what I have called *options*. But, even here, the modal element of the present view is quite different from that of previous accounts. It may be helpful to distinguish, here, between *weakly modal* and *strongly modal* phenomena. A weakly modal phenomenon is one such that true claims about that phenomenon *entail* claim about what is true at some possible world (or set of worlds) or other. A strongly modal phenomenon is one such that true claims about that phenomenon are *analyzable* in terms of what is true at some possible world (or set of worlds) or other. The distinction is crucial to appreciating, for instance, Kit Fine’s influential work on essence (Fine 1994). Fine argues that essence has a modal entailment: if something has a property essentially, then it has that property at every possible world. But Fine has argued influentially that the converse does not hold, and therefore that we should not give a modal analysis of essence. In the terms of the present discussion, Fine’s argument is an argument that essence is weakly, but not strongly, modal.

On the present view, ability too is weakly, but not strongly, modal. Claims about ability are analyzable in terms of genericity and claims about options. But options themselves are weakly, but not strongly, modal. In virtue of obeying the principle **Possibility**, they are weakly modal; but, as was emphasized in the discussion of that principle, they are not strongly modal. This was not an incidental feature of the account: it was precisely the rejection of a modal analysis of options that allowed us to explain why options, as well as abilities, do not satisfy the logical principle **K**.¹⁴

References

- Austin, J. L. (1956). Ifs and cans. In *Philosophical papers* (1961). Oxford: Clarendon Press
- Belnap, N., & Perloff, M. (1988). Seeing to it that: A canonical form for agentives. *Theoria*, 54, 175–199.
- Bhatt, R. (1999). Covert modality in non-finite contexts. *Dissertation*, University of Pennsylvania.
- Brown, M. (1988). On the logic of ability. *Journal of Philosophical Logic*, 17, 1–26.
- Carlson, G., & Pelletier, F. (Eds.). (1995). *The generic book*. Chicago: University of Chicago P.
- Cross, C. (1986). ‘Can’ and the logic of ability. *Philosophical Studies*, 50, 53–64.
- Davidson, D. (1967). The logical form of action sentences. In N. Rescher (Ed.), *The logic of decision and action*. Pittsburgh: University of Pittsburgh Press.
- Fara, M. (2010). Knowability and the capacity to know. *Synthese*, 2010, 53–73.

¹⁴ An earlier version of this paper was presented at the University of California, Riverside; I am grateful to the audience for discussion. Thanks also to Matthew Mandelkern and Michael Nelson for discussion and correspondence.

- Fine, K. (1994). Essence and modality. *Philosophical Perspectives*, 8, 1–16.
- Horty, J. (2001). *Agency and deontic logic*. Oxford: Oxford University Press.
- Horty, J., & Belnap, N. (1995). The deliberative stit: A study of action, omission, ability, and obligation. *Journal of Philosophical Logic*, 24, 583–644.
- Horty, J., & Pacuit, E. (Unpublished ms). Action Types in Stit Semantics.
- Kenny, A. (1975). *Will, freedom, and power*. Oxford: Blackwell.
- Kratzer, A. (1981). The notional category of modality. In H. J. Eikmeyer & H. Rieser (Eds.), *Words, worlds, and contexts*. Berlin: de Gruyter.
- Kripke, S. (1963). Semantical considerations on modal logic. *Acta Philosophica Fennica*, 16, 83–94.
- Maier, J. (2015). The agentive modalities. *Philosophy and Phenomenological Research*, 90, 113–134.
- Mandelkern, M., Schultheis, G., & Boylan, D. (ms). Agentive modals. *The Philosophical Review*. **(Forthcoming)**.
- Portner, P. (2009). *Modality*. Oxford: Oxford University Press.
- Spencer, J. (ms). Able to do the impossible. *Mind*. **(Forthcoming)**.
- Stalnaker, R. (1968). A theory of conditionals. In N. Rescher (Ed.), *Studies in logical theory*. Oxford: Blackwell.
- Thomason, R. (Unpublished ms). Defeasibly successful action.