Chance, Ability, and Control*

Matthew Mandelkern
New York University (mandelkern@nyu.edu)

May 23, 2023

Abstract

What does it take for S to be able to do φ? According to one analysis, it is enough that doing φ is compatible with S's intrinsic properties and contextually salient local circumstances. According to a competitor analysis, it is rather that S would do φ if she tried to. According to yet another alternative, it is that doing φ is somehow under S's control. I argue that probability judgments about chancy ability ascriptions provide an important new source of evidence for this debate. In particular, I argue that they provide evidence against the prima facie intuitive view that ability requires control over the action in question, and evidence for the view that ability ascriptions have an essentially conditional meaning.

1 Introduction

What does it take to be able to do something—say, wash the dishes before bed, read a paper, or hit a bullseye? According to a standard treatment, an ability ascription like 'I can hit a bullseye' is an existential modal claim, quantifying over possible worlds that hold fixed my intrinsic features and (contextually salient) extrinsic circumstances. So this sentence is true just in case, at some such world, I hit a bullseye.

According to a different idea (from Hume 1748), ability ascriptions encode a conditional fact: 'I'm able to hit a bullseye' is true just in case, if I try to hit a bullseye, I succeed. This *conditional analysis* of ability is big if true. It ties together the analysis of ability with the theory of conditionals, as well as with the philosophically important (if somewhat obscure) notion of trying.[1] It puts paid to a substantial research program in the semantics of modality in natural language, growing out of Kratzer 1977, 1981, which aims to develop a *uniform*

*I'm grateful to audiences at the University of Lisbon, the Centre for Human Abilities, the Dianoia Institute for Philosophy, and Johns Hopkins, and to David Boylan, Cian Dorr, Melissa Fusco, Ben Holguín, Bruno Jacinto, Joshua Knobe, Harvey Lederman, Annina Loets, Guillermo del Pinal, Daniel Rothschild, Paolo Santorio, Ricardo Santos, Ginger Schultheis, Malte Willer, and Snow Zhang for very helpful discussion.

[1] See Holguín and Lederman 2022 and citations therein for recent work on trying.

analysis of modal words, according to which all modals are existential or universal quantifiers over accessible possible worlds, with differences in their interpretation arising from differences in context. Most importantly, it provides an illuminating analysis of the words 'able' and 'can' (on their agentive uses), which, as Austin (1961) puts it, 'in philosophy... we seem so often to uncover, just when we had thought some problem settled, grinning residually up at us like the frog at the bottom of the beer mug'.

There are some obvious problems with the conditional analysis, as I explain in §2—reasons which have led many to reject it. However, as I will explain, recent work has tried to solve some of those problems with an approach that preserves the basic insight of the conditional analysis. But a deep objection remains open, based on the intuition that ability requires *control* (§3). If you do φ but only in an accidental or chancy way, does that show that you are *able* to do φ? Or does ability require something more substantial, like the possibility of doing φ in a controlled way? If the latter, then the conditional analysis is plausibly untenable, since the relevant conditional does not, on the face of it, encode any kind of control.

In this paper, I will argue that ability does not entail control. I will do so by introducing a new kind of evidence into the debate: namely, judgments about the *probabilities* of chancy ability ascriptions (§4). I will argue that these judgments show that ability does not require control (§5), and tell strongly in favor of some form of conditional analysis (§6) and, correspondingly, against some prima facie compelling objections to the conditional analysis (§7–§8).

2 The existential and conditional analyses

I'll start with a quick overview of some of the debate so far.

First, some preliminaries. My topic is *agentive modals*: words like 'able' and 'can' in English, on a reading where they are used to talk about abilities or their lack. (As I discuss in §8, in some cases it is unclear whether they are getting such a reading rather than a circumstantial reading. But for the most part I will focus on what I take to be paradigm cases of ability ascriptions. I move freely between 'able' and 'can'.) As is standard, I model agentive modals as denoting a relation between an individual and an action. I write $A_s\varphi$ for ⌜S is able to φ⌝ (on the relevant, agentive reading), and $\varphi(S)$ for ⌜S φ's⌝. I will be sloppy about use and mention (so I will both use φ as a variable over predicates in our target fragment, and as a metalanguage variable over actions).

2.1 The existential analysis

The natural starting point for a theory of agentive modals says that $A_s\varphi$ is an existential quantifier over accessible worlds. This assimilates 'able' to the $\Diamond$ of modal logic, and hence

to standard treatments of epistemic and circumstantial modalities. Analyses along these lines were proposed in Hilpinen 1969; Lewis 1976a; Kratzer 1977, 1981. While there are differences in implementation,[2] the basic idea is this:

Existential analysis: $[\![A_s\varphi]\!]^{c,w} = 1$ iff $\exists w' \in R_c(w) : [\![\varphi(S)]\!]^{c,w'} = 1$

Where ψ is a sentence, $[\![\psi]\!]^{c,w}$ is the truth-value of ψ at a world w and context c. R_c is the binary accessibility relation on worlds associated with c. So this semantic clause says that $A_s\varphi$ is true at a context c and world w just in case $\varphi(S)$ is true in some world which is accessible from w according to the context's accessibility relation. The idea is that the accessibility relation holds fixed salient facts about the agent's circumstance and her intrinsic features (see Vetter 2013 for a helpful characterization of the view), so that $A_s\varphi$ says that some world compatible with S's circumstances and intrinsic properties is one where she φ's.

So, for instance, a sentence like (1) is predicted to be true just in case Flo's circumstances and intrinsic features are compatible with her flying:

(1) Flo is able to fly.

If Flo is a penguin, (1) thus comes out false. If Flo is a swallow, and otherwise unimpeded from flying, then (1) comes out true. At first blush, this looks like a good theory.

However, other cases suggest that the existential account is too weak: it makes it too easy for ability statements to come out true. Consider this case from Mandelkern et al. 2017. Jo is playing darts. Jo's young daughter Susie exclaims:

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

Now suppose that Susie is an ordinary five-year-old child: she is relatively weak and uncoordinated, and it is extremely unlikely that she'll hit the bullseye if she tries. But it's not *impossible*. To make this more concrete, suppose that once Susie took a lucky shot and in fact hit a bullseye. So we know that it's *possible* for her to hit a bullseye. Still, most people won't readily assert or assent to (2). Intuitions about the precise status of (2) vary, but no one seems to think that (2) is clearly true. Instead, people tend to think that (2) is indeterminate, or false, or unlikely, or perhaps unassertable for yet some other reason. One of the goals of this paper is to clarify the precise status of sentences like (2). But the present point is that all of these judgments are, on the face of it, inconsistent with the existential theory, which predicts that (2) is clearly, determinately, certainly true. Why? Because it is clearly, determinately, certainly *compatible* with Susie's intrinsic features, and the present circumstances, that Susie

[2] In particular, Kratzer's treatment involves two contextual parameters, a modal base and ordering source, rather than one. However, in a fragment without conditionals, that treatment is equivalent to the present one, so this simplification should not matter for our purposes.

hit the bullseye on this throw. But that is just what the existential analysis says it takes for (2) to be true.

2.2 The conditional analysis

A natural first impulse in response to cases like (2) is to reach for some kind of *threshold* view, where abilities require that their prejacent not only be *compatible* with local circumstances and facts about the agent, but also be suitably *likely* or *normal*. However, a little reflection suggests that a theory like that wouldn't work, for we are able to do things that we are very unlikely to do, and which would be very abnormal. For instance, take Sylvia, a professor of philosophy with ordinary physical abilities and ordinary adherence to social norms who is in the midst of giving a colloquium talk. (3) is clearly true in this circumstance:

(3) Sylvia is able to remove her shoes and leave the room.

But it is extremely unlikely that she will do so, and it would be extremely abnormal for her to do so. So thresholds do not, on reflection, yield an easy fix to the existential analysis. We need some other way of strengthening the existential analysis.

A natural second response to these facts is to upgrade 'able' from an existential to a *universal* modal (that is, the $\Box$ of modal logic). This would make (2) clearly false. While this response has been mooted and, in some cases, advocated in different forms (e.g. Kenny 1976; Brown 1988; Giannakidou 2001; Giannakidou and Staraki 2012), on the simplest implementation it is clearly wrong. Sylvia is *able to* take her shoes off and leave right now, but there is no sense in which that is a necessity for her: some, indeed most, accessible worlds, in any reasonable sense of accessibility, are ones where she doesn't do so.

A much more promising idea treats ability ascriptions as involving an underlying *conditional* meaning. This is the conditional analysis, which was first put forward in Hume 1748 and has been an influential contender since (e.g. Moore 1912; Lehrer 1976; Cross 1986; Thomason 2005. On this account, ⌜S is able to φ⌝ says that if S *tries* to φ, then S does φ. That is, where $try(S, \varphi)$ is shorthand for ⌜S tries to φ⌝ and $>$ is the conditional operator ⌜If…then…⌝:

Conditional Analysis: $[\![A_s\varphi]\!]^{c,w} = 1$ iff $[\![try(S,\varphi) > \varphi(S)]\!]^{c,w} = 1$

To see the attractions of the account, consider the three examples we've looked at so far. Start with (2), 'Susie is able to hit the bullseye on this throw'. This has the same status in this theory as the corresponding conditional 'If Susie tries to hit the bullseye, she'll succeed'. Different theories of the conditional have different takes on the status of this conditional, but none predict that it is clearly true, as desired. According to Stalnaker 1968, 1981, this conditional is indeterminate in truth value. According to van Fraassen 1976, the conditional is very unlikely to be true. According to Adams 1975, the conditional has a low degree of

assertability. According to Lewis 1973; Kratzer 1981, the conditional is simply false. Any of these verdicts suffices to capture the basic intuition about (2): that you wouldn't naturally assert it or assent to it. More generally, it just seems intuitive that (2) is true just in case, if Susie tries to hit the bullseye, she succeeds. Next consider 'Flo is able to fly.' According to the conditional analysis, this is true just in case if Flo tries to fly, she succeeds. If Flo is a penguin, this latter conditional is clearly false; if Flo is an (otherwise unhindered) swallow, it is clearly true. Again, this exactly matches intuition about whether Flo is able to fly. Finally, consider 'Sylvia is able to remove her shoes and leave the room.' This is true, on the conditional analysis, just in case if Sylvia tries to remove her shoes and leave, she succeeds. But that is clearly true, and so the ability ascription is rightly predicted to be true.

2.3 Problems for the conditional analysis

So far, so good. But the conditional analysis as stated just now has decisive problems. A recent paper, Mandelkern et al. 2017, aimed to solve at least some of those problems with a variant of the conditional analysis. In this paper, I want to grant that that variant, or some view like it, can indeed solve those problems, and focus on a different, and potentially deeper, problem with the conditional analysis, concerning the relationship between ability and control. In the rest of this section, I'll very briefly describe the problems which Mandelkern et al. (2017) aim to solve, and explain their solution. Since my broader goal is to set these problems aside, readers more interested in the big picture than the details might want to skip to the beginning of the next section.

Mandelkern et al. (2017) focus on three problems for the conditional analysis. The first concerns *negated* ability statements. According to the conditional analysis, ⌜S is not able to φ⌝—that is, $\neg A_s\varphi$—is true just in case it's *not* true that S will φ if she tries to φ. But note that this doesn't say anything about what happens when S *doesn't* try to φ. In fact, according to the conditional analysis, the truth of ⌜S is not able to φ⌝ is consistent with S *succeeding* at φ when she doesn't try to φ. That's intuitively wrong: ⌜S is not able to φ⌝ entails that there's *nothing* S can do that will result in doing φ, contra the conditional analysis. Concretely, consider Mirabel, whose visual representations are inverted across the vertical axis. She's an excellent dart player, so if she tries to hit the left side of the dartboard, she'll hit the right, and vice versa. Furthermore, Mirabel is unaware of this inversion. Now consider (4):

(4) Mirabel is unable to hit the left side of the dartboard.

Intuitively (4) is false (or at least has a prominent false reading), because Mirabel *can* hit the left side: all she has to do is try to hit the right side. But according to the conditional analysis, (4) is (clearly, unambiguously) true, because if Mirabel tries to hit the left side of the dartboard, she'll hit the right side.

The second problem comes from cases where the conditional analysis is too weak. For instance, to use a case from Thomason 2005, suppose John has made plans to go to a movie. Ann invites him to dinner and he replies:

(5) Sorry, I can't go, I'm going to a movie.

(5) appears perfectly assertable in this context. But the conditional analysis predicts that the negatum in (5), 'I can go', is *true*, since it is clearly, unambiguously true that, if John tries to go to dinner, he'll succeed.

The third problem arises from cases where the conditional analysis appears too strong. For instance, consider (6), based on Vranas 2010:[3]

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

(6) is intuitively true—David is normal, breathing-wise—but if David *tries* to breathe normally, he'll focus on breathing normally and then will fail to do so; so the conditional analysis predicts that (6) is false.

2.4 The Act Conditional Analysis

These problems show that the conditional analysis in its standard form cannot be right. But Mandelkern et al. (2017), developing an idea of Chisholm 1964's, argue that the spirit of the conditional analysis can be saved with a relatively minor revision: namely, by adding existential quantification over a contextually supplied set of actions. In particular, their view—the *act conditional analysis*, or *ACA*—says that S is able to φ just in case there is *some* contextually salient action ψ such that, if S tries to do ψ, she does φ. More carefully, where $\mathcal{A}_{c,s}$ is a set of actions which are in some sense "practically available" to S in context c:

$$\textit{Act Conditional Analysis}\text{: } [\![A_s\varphi]\!]^{c,w} = 1 \text{ iff } \exists\psi \in \mathcal{A}_{c,s} : [\![try(S,\psi) > \varphi(S)]\!]^{c,w} = 1$$

There is more to say about the details of this account, in particular the notion of practically available actions, but for the sake of brevity, let me just highlight how this approach can solve the three problems just sketched.

First, $\neg A_s\varphi$ on this account means that there is *no* contextually salient action ψ such that, if S tries to do ψ, she does φ. That means, for instance, that the ACA rightly predicts that 'Mirabel is unable to hit the left side of the dartboard' is false (provided *hitting the right side of the dartboard* is practically available), since there *is* something (hitting the right side) such that if Mirabel tries to do that, she hits the left side.[4]

[3] For a related case, consider the golfer in Austin 1961.

[4] This sentence plausibly also has a true reading, since Mirabel doesn't *know* to aim right. The ACA can predict this, too, in a context where the practically available actions are *hitting the right side, not hitting the right side*, since neither of these is such that if Mirabel tries to do it, she hits right.

Second, consider cases where the conditional analysis is too weak, like that of John, who says that he can't go to dinner because he's going to a movie. Of course, if he tries to go to dinner, he succeeds. But if we treat *meeting Ann for dinner* as practically unavailable to John because of his other commitments, then the ability ascription 'I'm able to go to dinner' won't come out true. In other words, if the only practically available actions in this context are ones compatible with John's existing commitments, then the ACA rightly predicts a true reading of 'I can't go to dinner', in line with intuitions.

Finally, consider cases where the conditional analysis is *too strong*, like that of David, who can breathe normally. How can he do this? Well, by trying to do something else, say, play piano for a few minutes. So the ACA rightly predicts there is a true reading of 'David can breathe normally', since there is *something* such that if he tries to do it, he breathes normally.

So the extra quantificational resources of the ACA let it avoid the problems we just surveyed for the simplest version of the conditional analysis. There is a cost, to be sure: introducing the set of practically available actions adds an extra degree of freedom to the analysis. But I want to grant that the ACA, or some view like it, can solve the three problems sketched above, so that some version of the conditional analysis remains a contender view.

3 Control and ability

In the rest of the paper, I will address a remaining, and I think potentially much deeper, question about the potential adequacy of *any* broadly conditional analysis, concerning the connection between agency and *control*.

The worry is this. If you try to φ and succeed, then according to the conditional analysis—and the ACA, and indeed any reasonable view in the neighborhood of the conditional analysis—it follows that you were able to do φ, since from the fact that you tried to do φ and succeeded, it follows (on any reasonable conditional logic) that if you tried to φ, you succeeded. This follows *even* if φ-ing is a total fluke, accidental, haphazard, out of your control. In the recent literature, several authors have objected that this neglects an important connection between agency and *control*: being able to do φ requires having φ somehow *under your control*, so that *flukily* doing φ doesn't necessarily suffice to be able to do φ.

To see the worry, return to Susie. She will wildly throw a dart at a dartboard, trying to hit the bullseye. There is some chance that she will hit the bullseye, just by luck, a random fluke. In that case, according to any reasonable version of the conditional analysis,[5] she is

[5]There are two ways to push back. First, we could adopt a theory of the conditional which invalidates Strong Centering, the principle that when ψ and χ are true, so is $\psi > \chi$. But Strong Centering is nearly universally accepted; nor would this fit with the intuitions about probability which I elicit presently. Second, in the framework of the ACA, we could argue that φ may be practically unavailable even if you try to do φ, but this would lead to a number of peculiarities—for instance, the coherence of 'Susie hit the bullseye but she was unable to'.

able to hit the bullseye, since she will have tried to hit the bullseye and succeeded. But since Susie is all over the place, many argued that this cannot be right: flukily hitting a bullseye does not suffice to be *able to* hit the bullseye. To be able to hit the bullseye, you have to do something more than just flukily hit it: the action of hitting the bullseye must somehow be *in your control*. And so the *Success* inference is not valid, pace any form of conditional analysis:[6]

Success: $try(S, \varphi) \wedge \varphi(S) \vDash A_s\varphi$

The intuition that ability requires control, and thus that Success is invalid, is widespread. Here are a few representative quotes in the recent literature, many of which arise in the context of pushing back against some form of conditional analysis:[7]

- 'ability requires control' (Loets and Zakkou, 2022);
- 'ability ascriptions [are] a kind of *hypothetical guarantee*. When someone says 'John can go swimming this evening', she is informing her interlocutors that going swimming this evening is, in a certain sense, within John's control' (Mandelkern et al., 2017);
- 'accidental, or *fluky*, success is insufficient for ascriptions of ability' (Fusco, 2020);
- 'control is central to ability. . . the claim that I can surf that wave is strong—it says that surfing that wave is within my control' (Boylan, 2020);
- 'the can of ability is essentially an existential quantifier over a set of available actions, and an action is available to an agent just in case he or she is deemed to have sufficient understanding of how to achieve the relevant outcome. . . [that is,] a good chance at succeeding in performing the relevant action, should he or she try to do it' (Willer, 2021).

For a more direct attack on *Success*—the antipode to the thesis that ability requires control—here is a discussion from Kenny 1976:

[6]The name is sometimes used for the stronger principle that $\varphi(S) \vDash A_s\varphi$. But this latter principle is not actually validated by the simple conditional analysis. It is validated by the ACA, provided that whatever the agent actually tries to do counts as practically available. But to keep the dialectic simple, I'll focus on the weaker principle in the text; I don't think a lot turns on this choice, since the arguments I give also speak in favor of the stronger principle.

[7]The dialectic in some cases is complicated. Loets and Zakkou (2022), while arguing that the control intuition has merit, are primarily concerned with bringing out a conflict between control and claims about the duals of agentive modals, rather than arguing for one resolution of that conflict. Mandelkern et al. (2017) endorse the control intuition but give a theory which does not really link ability to control, as we will see. Boylan (2020) develops a variant of the ACA which aims to capture the control intuition by introducing a distinction between future and past-oriented ability ascriptions, with Success valid for past-oriented ascriptions but not future-oriented ones, which are instead indeterminate. It is still possible that Boylan's view is consistent with the judgments I elicit below, however, if indeterminate sentences can have non-extremal probabilities.

> I cannot spell 'seize': I am never sure whether it is an exception to the rule about 'i' before 'e'; I just guess, and fifty times out of a hundred I get it right. On each such occasion we have a counter-example to [*Success*]: it is the case that I am spelling 'seize' correctly but it is not the case that I can spell 'seize' correctly.

The recent literature also furnishes a number of proposals which aim to capture the control intuition, and thus invalidate *Success*. To give a better sense of the control intuition, I'll give a brief and informal sampling of some of those proposals. First, we can encode control in the *truth-conditions* of ability ascriptions. This is roughly the proposal of Fusco 2020. Following the tradition of Brown 1988; Horty and Belnap 1995, Fusco treats ability ascriptions as complexes of existential and necessity operators: $A_s\varphi$ means, in short, that it is historically possible that S's powers *necessitate* $\varphi(S)$ (so $A_s\varphi$ is equivalent to $\Diamond\boxdot\varphi(S)$, where $\Diamond$ is a historical possibility operator and $\boxdot$ quantifies over worlds compatible with 'everything the agent's powers are able to necessitate'). While Fusco doesn't put things explicitly in terms of control, we can see a view like this as a way of encoding the control intuition, cashed out in terms of necessitation: $\ulcorner$S can $\varphi\urcorner$ says that there is a possibility where S does φ *in a controlled way*. Another approach is to encode the control intuition via a threshold: for instance, Willer (2021) suggests that for S to be able to φ is to have 'a good chance at succeeding in performing the relevant action, should he or she try to do it': that is, it requires that φ is in the agent's control to a sufficient degree, in the sense that trying to do φ results in performing φ enough of the time (cf. Jaster 2020). A final approach is to encode control as a *presupposition* of ability ascriptions. This idea is inspired by a recent proposal of Santorio (2022). On Santorio's account, $A_s\varphi$ says that it is possible that S does φ, and presupposes that S has a state which is causally sufficient for φ in any accessible possibility where S in fact does φ. Causal sufficiency is, in turn, a necessity-like notion, spelled out in terms of causal models. While Santorio doesn't gloss causal sufficiency in terms of control, it is natural to see it (like Fusco's notion of necessitation) as a generalization of the notion of control, since, among other things, it is intended to rule out ability ascriptions in cases like that of the haphazard but lucky dart player.

To situate these views in the context of our running example: I take it that hitting the bullseye is not in Susie's control, since it is enormously unlikely that she will hit it when she tries; if she does hit it, we would describe her hitting it as a matter of luck and chance. (If an action like this were taken to be in Susie's control, then the control intuition wouldn't have any bite at all. At least some of the authors just cited discuss cases like that of Susie, explicitly noting that in a case like this, the relevant control/necessity/sufficiency condition is not met.) So all these views predict that it is *clearly not true* that Susie is able to hit the bullseye. By contrast with all these views, since it is *possible* that Susie will hit the bullseye, it follows on the conditional analysis that it is *possible* that she will be able to hit the bullseye.

It should be noted that, as Mandelkern et al. (2017) discuss, the conditional analysis can account for part of the control intuition by adverting to *genericity*. Consider the *general* ability ascription in (7):

(7) Susie can hit bullseyes.

(7) intuitively says: Susie is *generally* able to hit a bullseye. Hitting a bullseye just once, by accident, does not show that Susie is generally able to hit a bullseye, even on a view where Success is valid. So, even though the conditional analysis validates Success, it doesn't validate the inference from ⌜S does φ at t⌝ to ⌜S is able to do φ in general⌝ or ⌜S (generally speaking) has the ability to φ⌝. In other words, on (any form of) the conditional analysis, the fact that Susie will try and succeed at hitting a bullseye at 3 pm entails the truth of (8), but not of the generic (7).

(8) Susie can hit a bullseye at 3 pm.

However, it is not clear that this specific vs. generic distinction suffices to capture the control intuition; indeed, the proposals cited above all maintain that it is not just *generic* ability ascriptions which require control, but also specific ones like (8) (that is, ability ascriptions whose complement is a fully tensed action). So in what follows, I'll continue to focus on specific ability ascriptions, which is where the heart of the controversy about control lies.

The classic argument that ability requires control comes from Kenny 1976.[8] Suppose Alice shuffles a standard deck of cards and places it face down. At 3 pm she will draw a card at random from the deck. Consider (10-a) and (10-b):

(10) a. Alice can draw a red card at 3 pm.
b. Alice can draw a black card at 3 pm.

According to Kenny 1976, both (10-a) and (10-b) are false. Since Alice doesn't have *control* over the color of the card she draws, she is neither able to draw a red card nor able to draw a black card.

But note that Alice *will* draw a red card or a black card. If doing φ entails being *able* to do φ, then it follows that either she can draw a red card, or she can draw a black card,

[8] Another argument, from Santorio 2022, comes from conditionals like (9):

(9) If Susie hits the target out of sheer luck on this throw, then Susie is able to hit the target on this throw.

Santorio argues that a conditional like (9) does not seem like a logical truth, but it should if Success were valid. (9) is certainly an odd sentence to produce, but so are many other logical truths, and it's not clear exactly what intuitions here are tracking; the balance of evidence against the control intuition makes me somewhat inclined to think that (9) just is a logical truth, after all.

contrary to Kenny's intuitions.[9]

For another case in a similar vein, consider this variant from Boylan 2020:

> I am a fairly bad dart player. I regularly hit the bottom half when I aim for the top, and vice versa. But I never miss the board entirely. I am about to take a shot. I am skilled enough to know I will hit the board; so I know the following:
>
> (11) I will hit the top half of the board on this throw or I will hit the bottom half of the board on this throw.
>
> But it does not seem that I should ascribe myself either of the following abilities here:
>
> (12) a. I can hit the top on this throw.
> b. I can hit the bottom on this throw.
>
> Even their disjunction does not seem true.

Once again, if Boylan's intuitions here are correct, it is because doing φ in an uncontrolled, fluky way does not suffice to be *able* to do φ: that is, because ability requires control.

4 Chancy abilities

I feel the pull of these intuitions. They pose a compelling challenge to Success, and hence to any form of conditional analysis. But I've become convinced that ability does not require control: I think success, no matter how fluky, entails ability. What convinced me was probability judgments about ability ascriptions, which I think provide very strong evidence for Success and against the control intuition. I will begin in this section by eliciting intuitions about chancy abilities, before going on in the following sections to explore their ramifications.

Start by focusing again on Susie. She is a haphazard dart player, tossing darts at he board; she can barely hit the dartboard, let alone the bullseye. But every once in a while, she gets a bullseye, just by chance; say this happens once every thousand throws or so. So the probability that she'll hit a bullseye on any particular dart throw is about .1%. (It doesn't matter exactly what sense of probability we have in mind in these cases. I will move freely between talk of chance and probability, and between talking about the probability of sentences and of the

[9]Note that, in order to fit the schema in Success, it must be that either Alice will *try* to draw a red card at 3 pm and succeed, or she'll *try* to draw a black card at 3 pm and succeed. Actually, we don't know that one of those things will happen in this case. But I want to set aside this response to the argument, because the ACA predicts not only that the inference from $try(S, \varphi) \wedge \varphi(S)$ to $A_s\varphi$ is valid, but also that the inference from $\varphi(S)$ to $A_s\varphi$ is valid, provided that whatever S in fact tried to do is treated as practically available. I think there are reasons to follow the ACA here, so Kenny's argument still targets the most plausible extant version of the conditional analysis.

corresponding propositions.) Now suppose that when the clock strikes 3 pm, Susie will throw the dart at the dartboard. Consider the question in (13):

(13) What's the chance that Susie will be able to hit a bullseye at 3 pm?

The most natural answer to (13), I maintain, is .1%: that is, the probability that Susie will be *able* to hit a bullseye is just the probability that Susie *will* hit a bullseye, *conditional* on trying to.

An immediate worry about this case is that .1% is *also* the probability that Susie *will* hit a bullseye at 3 pm, since we are sure she will try to. So an error theory about this judgment is that we might simply be answering a different question in the neighborhood, namely, about the probability that Susie *will* hit a bullseye, not that she will be *able* to. However, this worry can be put to rest by considering a variant on the case where you're not sure whether Susie will take a shot at 3 pm; say there is a 50% chance she will, and a 50% chance she won't. Given that, the chance that Susie *will* hit a bullseye at 3 pm is .05%. But the chance that she *can* hit a bullseye intuitively remains .1%: that is, it remains the chance that she will hit a bullseye, *conditional on trying to*.[10]

For variety, I'll give a few cases with a similar structure. Suppose next that Ludwig is going to an audition. Consider (16):

(16) What's the probability that Ludwig can play the Hammerklavier sonata through at the audition without making an error?

[10] A related worry is about *actuality entailments*: past-oriented ability ascriptions have a prominent interpretation where they are equivalent to their prejacents, so, e.g., 'Susie was able to get the groceries' has a prominent reading where it feels equivalent to 'Susie got the groceries' (see e.g. Bhatt 1999; Hacquard 2010). I am sidestepping this issue by focusing exclusively on present and future oriented ability ascriptions. Of course, you might be worried about actuality entailments there too, but actuality entailments in the future oriented case would not explain credence judgments, since the credences about abilities elicited here *don't* match the probabilities of the prejacents, as they would if actuality entailments were coloring judgments. To be sure, out of the blue it is very natural to hear (14) and (15) as equivalent, and to assign them the same probability:

(14) Annina can finish the proofs by tomorrow.

(15) Annina will finish the proofs tomorrow.

But this could simply be because we naturally assume, in evaluating (14), that Annina will try. And indeed, when we make explicit that we are not assuming this, probability judgments about (14) and (15) do diverge. So, if we have .5 credence that Annina will try, and credence .9 that she'll finish if she tries, we should intuitively have credence .9 in (14) and credence .45 in (15). This suggests that, while there are natural contexts where we interpret ability ascriptions in the same way as their bare prejacents, this is not what is responsible for the probability judgments I am eliciting.

A related question is about conditionals like 'If Annina tries, she'll be able to finish the proofs'. On a conditional analysis, there is something redundant about such a conditional, since it means the same thing as 'Annina will be able to finish the proofs'. I suspect that adding the redundant conditional antecedent is a way of bringing out that the claim in question is a genuine ability ascription, not a roundabout claim that she *will* finish the proofs.

Suppose your credence that Ludwig *will* play the sonata through without making an error, conditional on him trying, is .2. Then, intuitively, your answer to (16) should be .2. Once again, we can ensure that you aren't just targeting the prejacent by making it an open question whether Ludwig will be asked to play the Hammerklavier or something else. This should lower your credence in the proposition that Ludwig *will* play the Hammerklavier without an error. But intuitively it doesn't change judgments about (16).

Or consider Ginger, who is standing on the basketball court, considering whether to attempt a free throw. To make things concrete, suppose that Ginger is 50% likely to take a shot. And conditional on taking the shot, she is 10% likely to make a basket (she's taken hundreds of free throws over the last few weeks, and made 10% of them). What's the chance of (17)?

(17) Ginger can make this shot.

Intuitively, 10%. Again, this matches the chance of Ginger making the shot, conditional on trying. And again, it does *not* match the chance of Ginger making the shot simpliciter (which is only 5%).

For a final case, consider Benjy, an otherwise very good cat who really doesn't like getting into his carrier for vet visits. Based on past experience, I have about a 20% rate of success at getting him into his carrier. Given that, what is the chance of (18)?

(18) I can get Benjy into his carrier for this vet visit.

Intuitively, 20%. Again, this matches the chance of success conditional on trying. Once again, we can pull apart this judgment from the chance of getting Benjy into his carrier, by making it chancy whether I take Benjy or his sister Little Cow to the vet: I'll try to take whichever cat I see first. So the chance that I actually get Benjy into his carrier is much lower than 20%.

These are my intuitions, anyway, and match my informal polling.

5 Against control

I'll now turn to the significance of these judgments. I'll start with a negative claim: these judgments show that ability does not require control.

For concreteness, I'll focus on the first case, involving Susie, who is haphazardly chucking darts at the dartboard. Since she has a .1% chance of getting a bullseye on any given throw, there's intuitively a .1% chance that she *can* hit the bullseye at 3 pm. This is a problem for analyses which tie ability intimately to control. I'll assume that Susie doesn't have control over the action of hitting a bullseye, in the sense relevant to ability. Again, if the control intuition is meant to have any bite, it rules out ability in cases like that of Susie.[11]

[11] If you think that a .1% chance of success is enough for control, lower the rate as much as you like. For any

So what do analyses that tie ability to control predict in a case like Susie's, specifically about the answer to the question: what is the probability of (19)?

(19) Susie will be able to hit a bullseye at 3 pm.

It depends a bit on how exactly control is incorporated. Above I briefly surveyed three approaches; I'll go through their verdicts about (19) in turn. On the first approach, exemplified by Fusco 2020, a control requirement is part of the truth-conditions of ability ascriptions: $A_s\varphi$ is true just in case there is some historically possible world where S's powers *necessitate* $\varphi(S)$. Is there an accessible world where Susie's powers necessitate her hitting a bullseye? Of course, powers necessitating something is a pretty theoretical notion, so it is not easy to intuit the answer to this question. But, again, this clause is supposed to capture the connection between ability and control, and Susie doesn't have control over hitting a bullseye. So on this account, (19) is predicted to be false. Not only that, but we should be *sure* it's false: we're sure that Susie doesn't have the relevant kind of control over hitting the dartboard. So, on a view like Fusco's, the probability of (19) is 0: there's *no* chance that Susie will be able to hit a bullseye at 3 pm, because there's no chance that her powers necessitate hitting a bullseye. But this is clearly the wrong verdict. There's *some* chance that Susie will be able to hit a bullseye; not a lot, but some. And that's enough to show that ability doesn't require control in the straightforward, truth-conditional way that Fusco's accounts encodes.

(To be sure: sometimes we can assert things that we aren't sure of, if we're very confident of them. I can tell you that my car is parked two blocks away, even if I have only, say, .99 credence in that. Likewise, in the case of an improbable ability like Susie's, I might tell you that Susie won't be able to hit a bullseye, since I have very high credence—.999—that she won't be able to. But just because I can reasonably say it doesn't mean it's true. An account like Fusco's predicts not only that 'Susie won't be able to hit the dartboard' is assertable, but also that you should be *sure* it's true; and that just seems clearly false here.)

What about a more roundabout connection between ability and control? We briefly surveyed two other options above. One of those, suggested by Willer (2021)'s informal remarks, appeals to a notion of a *threshold*: to be able to do φ, S must have a 'a good chance at succeeding in performing the relevant action, should he or she try to do it'. Well, Susie doesn't have that. Of course, there is flexibility in what counts as a good enough chance in this definition; but, again, if the notion is supposed to do any work, this is a threshold that someone like Susie clearly doesn't pass. So, again, the prediction is that we should be sure that (19) is false, because we are sure that the corresponding threshold judgment in (20) is false.

ϵ, no matter how small, if Susie has an ϵ chance of hitting a bullseye when she tries, then, intuitively, she has at least an ϵ chance of being able to a hit a bullseye. Surely there is an ϵ low enough that Susie does not have control over hitting a bullseye. But then all the present points can be made with that version of the case.

(20) There is a good chance that Susie will hit the dartboard at 3 pm if she tries.

So a threshold analysis should predict, again, that we judge (19) to have no chance at all of being true, just as we judge (20) to have no chance at all of being true. But, again, this is apparently wrong: we should be sure that (20) is false, but we should assign non-zero credence to (19).

The last approach, from Santorio 2022, says that ability and something in the neighborhood of control are connected via a *presupposition*: on Santorio's view, $A_s\varphi$ says that $\varphi(S)$ is possible, and also presupposes that in any possibility where S does φ, she has a state which is causally sufficient to bring about φ. As Santorio discusses, an agent like Susie doesn't meet this condition: the circumstantial possibilities where she hits a bullseye are ones where she does so haphazardly, by chance, not thanks to causal necessity. So, on Santorio's account, we are *sure* that (19) has a false presupposition. Usually when we are sure that a sentence has a false presupposition and we are asked to judge its probability, we find the question ill-formed, and, at best, can get a judgment of 0 (a judgment that arises when we 'locally accommodate' the presupposition and interpret it as part of the asserted content), as in sentences like (21) (where the target proposition presupposes that Liam has missed a rent payment in the past) or (22) (where the target proposition presupposes that Alyssa once drank):

(21) Liam has never missed a rent payment. What's the chance that he'll miss another one?

(22) Alyssa has never touched alcohol. What's the chance that she quit drinking?

In light of judgments like this, a presuppositional view predicts that a question like 'What is the chance that Susie will be able to hit a bullseye at 3 pm?' will strike us as ill-formed, since we know it has a false presupposition. At best, a view like this predicts that, if we are forced to answer the question, we will be able to access a judgment of 0.

Note further that even if we maintained that for some reason we are able to ignore the presupposition in this case, that would yield not the observed judgment of .1, but rather a judgment of 1 (since we are *sure* that there is a circumstantially accessible world where Susie hits a bullseye, and that is all the ability ascription requires for truth once we ignore its presupposition). It's hard to see a route for an approach like this to make sense of the observed judgments.

While there might be other ways of connecting ability ascriptions to control (or some control-like notion) beyond the three I have sketched here, I suspect that all of them will run aground on intuitions about chancy abilities. If ability requires control, then we can be sure that one of the requirements for (19) to be true is not met, and hence that it does not have any chance at all of being true. But (19) manifestly does have some small chance of being

true.

I have focused on the case of Susie here for concreteness; the same points can be made with any of the other scenarios I considered in the last section.

6 In favor of Success

So much for the negative lessons of probability judgments about ability ascriptions. What about positive lessons? In this section I'll argue that probability judgments speak in favor of the validity of Success, and moreover support some form of a conditional analysis of ability.

Start with Success, the inference pattern from $try(S, \varphi) \wedge \varphi(S)$ to $A_s\varphi$. One way of summarizing the probability judgments elicited in §4 is this: the probability of $A_s\varphi$ is always at least as great as the probability of $try(S, \varphi) \wedge \varphi(S)$. Hence, for instance, your credence that Susie will be *able* to hit a bullseye should be at least as great as your credence that she will *try and succeed*. In general, it is a law of probability that, when ρ entails χ, the probability of χ is always at least as great as the probability of ρ. So the validity of Success would provide a *partial explanation* of the patterns brought out in §4. This provides a powerful (albeit indirect) argument in favor of Success—and hence in favor of accounts that validate Success.

The probability judgments elicited in §4 also provide a more direct argument in favor of some form of a conditional analysis. The observation is that, in all the cases we looked at there, intuitive judgments about the probability of $A_s\varphi$ exactly match intuitive judgments about the corresponding conditional $try(A, \varphi) > \varphi(S)$. In other words, the following pairs all intuitively have the same probabilities:

(23) a. Susie will be able to hit a bullseye at 3 pm.
b. If Susie tries to hit a bullseye at 3 pm, she'll succeed.

(24) a. Ludwig can play the Hammerklavier sonata through without making an error.
b. If Ludwig tries to play the Hammerklavier sonata through without making an error, he'll succeed.

(25) a. Ginger will be able to make this shot.
b. If Ginger tries to make this shot, she'll succeed.

(26) a. I can get Benjy into his carrier for this vet visit.
b. If I try to get Benjy into his carrier for this vet visit, I'll succeed.

Intuitively, for instance, (23-a) has probability .1%, and *so does* (23-b): they both have the probability of Susie hitting a bullseye conditional on trying. This provides a powerful argument for a theory on which the sentences above are (at least to a first approximation) pairwise equivalent.

In more detail: the probability judgments elicited in §4 tend to match the *conditional probabilities* of doing the relevant action, conditional on trying. These probabilities are also, intuitively, the probabilities of the corresponding *try*-conditionals $try(S, \varphi) > \varphi(S)$. The relation between conditional probabilities and probabilities of conditionals is famously vexed (see Stalnaker and Thomason 1970; Adams 1975; Lewis 1976b; see Khoo and Santorio 2018 for an excellent recent overview). But nearly everyone agrees that, for simple conditionals (conditionals which don't embed modals or conditionals), there is at least a prominent default interpretation on which their probabilities equal the corresponding conditional probabilities of their consequents conditional on their antecedents. Hence to give just one example, consider (27):

(27) If I flip the coin, it will land heads.

It seems that the probability of (27) should just match the probability of *heads* conditional on *flip* (.5 if the coin is fair, higher if it is biased towards heads)

So given an approximate connection between the probabilities of conditionals and corresponding conditional probabilities, the judgments elicited in §4—which matched the probability of success, conditional on trying—tell in favor of some form of the conditional analysis on which the pairs in (23)–(26) mean roughly the same thing.

Of course, all the reasons that we saw for rejecting the simple conditional analysis still apply when we turn to probability judgments. If you know John has made plans to see a movie, then you know that there is *no* chance that he can go to dinner with Ann, even while you would judge that the corresponding conditional, 'If John tries to go to dinner, he'll succeed' is very likely. So cases like this (still) motivate a move away from the simplest version of the conditional analysis to a variant like the ACA. But the ACA approximates the simple conditional analysis in many cases—in particular, if φ and $\overline{\varphi}$ are the practically available actions, and trying to do $\overline{\varphi}$ won't result in doing φ, then the ACA just collapses into the simple conditional analysis. So probability judgments tell in favor of *some* form of conditional analysis where the pairs in (23)–(26) mean roughly the same thing; this could be the ACA, or some other variant of the simple conditional analysis. (A further question is whether the ACA gets things right in cases where it doesn't collapse into the simple conditional analysis. I won't explore this here, since I am not engaging in this paper in intra-party debates about the correct form of the conditional analysis, but instead trying to motivate adopting *some* form of conditional analysis.)

A natural thing to ask for at this point is a semantic model for the probabilities of conditionals which, together with some form of the conditional analysis, would yield all the judgments we've seen so far. Easy: just pick your favorite model for the probabilities of conditionals—there are a number of viable contenders (for a sample, consider any of

van Fraassen 1976; McGee 1989; Kaufmann 2009; Bacon 2015; Goldstein and Santorio 2021; Khoo 2022)—and combine it with your preferred version of the conditional analysis, and you'll have a model for the judgments elicited so far.

A dialectical subtlety: there are analyses of the conditional on which it encodes a kind of necessity. If we coupled a conditional analysis of ability with a necessity analysis of the conditional, then we would link ability to control, after all. So not *any* conditional analysis can make sense of the judgments I've elicited here. For instance, on Lewis (1973); Kratzer (1981)'s influential theories, a conditional is a restricted *necessity* operator, which says that its consequent is true at *every* "closest" antecedent world. If you adopted a theory of the conditional like that, together with some form of the conditional analysis of ability, you'd get a theory that vindicates the control intuition—and which doesn't conform to probability judgments. For instance, on Kratzer's theory, the conditional 'If Susie tries to hit the bullseye, she'll succeed' is a restricted 'must' claim, equivalent to 'If Susie tries to hit the bullseye, she *must* succeed'. This is, intuitively, true just in case Susie has total control over hitting the bullseye. Thus in our case, it is *sure* to be false: even if she is to hit the dartboard by luck, it's certainly not the case that she *must* hit the bullseye. So if we adopted the conditional analysis but then analyzed the conditional in Kratzer's terms, we would wrongly predict that the chance that Susie can hit the bullseye is 0. So, more carefully, what probability judgments support is a form of conditional analysis which is in turn implemented with a theory of the conditional able to account for probability judgments about the corresponding conditionals elicited here. Again, there are many possible implementations that fit the bill,[12] and there is no need to commit to one here.

7 Kenny's argument

So probability judgments provide a powerful argument against incorporating control into our analysis of ability, and in favor of a broadly conditional analysis. But what should we say about Kenny's powerful argument for the control intuition? I will argue that probability judgments, plus some observations about peculiarities of that case, provide a way to defuse that argument.

Recall the case. Alice shuffles a deck of cards and places it face down. At 3 pm she will draw a card at random from the deck. It will be either red or black, so if success entails ability, either (10-a) or (10-b) is true:

(10-a) Alice can draw a red card at 3 pm.

(10-b) Alice can draw a black card at 3 pm.

[12] Rothschild (2013) has even argued that a necessity-based analysis of the conditional can capture judgments about conditional probabilities.

But according to Kenny, these are both intuitively false, since she lacks the relevant control.

Now what's clearly right in Kenny's case is that you shouldn't *say* either (10-a) or (10-b). But, of course, just because something isn't assertable, it doesn't follow that it's false: there are many things that make a sentence unassertable. There is something in particular that makes both (10-a) and (10-b) very strange, which is that neither of the relevant actions—*drawing a red card* and *drawing a black card*—would be reasonable things for Alice to *try* to do, since she should not think she can influence the color of the card. It looks to me like $A_s\varphi$ is only assertable, in general, when we leave it open that S could (in some broadly circumstantial sense) try to φ. And in this case, it's hard to imagine Alice *trying* to draw a red card. What would that be like, for her? She might scrunch up her mind and think hard about red. But most people would, I think, not even know what it would be to *try* for red in a case like this.

We can get around this by making a very simple change to the case: just give Alice the irrational belief that she might be able to influence the color of the card she picks, say by thinking hard about the relevant color, so that it is more plausible that she could *try* to draw a red card, and to try to draw a black card. The crucial thing is that this doesn't actually change the structure of the case vis-à-vis Alice's control over the outcome. She still doesn't have control in any sense over the color of the card she draws. But making these changes makes (10-a) and (10-b) easier to countenance. And, in particular, it brings out a very clear probability judgment: there is a 50% chance that Alice can draw a red card at 3 pm, and a 50% chance that she can draw a black card at 3 pm. (It doesn't matter exactly how you get here: change the case so that we can easily imagine that Alice might *try* to draw red, and might try to draw black, and these judgments seem to come clearly into view.)

But if that's right, then Kenny is wrong about this case. It's not that Alice is *unable* to draw a red card and unable to draw a black card. If ability required control, then that would be correct. But probability judgments (again) suggest that ability doesn't require control. Instead, they suggest that one of (10-a) or (10-b) is true—it just depends on what color card Alice in fact will draw. We don't know which it is, so we aren't in a position to assert either one of these; instead, we distribute our credence evenly over them.

Or recall Boylan's variation on Kenny's case: David is a bad dart player. He regularly hits the bottom half when he aims for the top, and vice versa. But he never misses the board entirely. Boylan claims that in this case, neither (28-a) nor (28-b) is true, even though David will hit the bottom or the top of the dartboard on this throw:

(28) a. David can hit the top on this throw.
b. David can hit the bottom on this throw.

Now, if I'm right that asserting ability ascriptions requires, in some sense, the circumstantial possibility of the agent in question trying the action in question, then we should make sure

that David could try to hit the top, and could try to hit the bottom, to avoid independent noise from this issue. So let's say that David always aims at a particular quadrant of the dartboard, in an effort to improve his throw. So on this throw, he'll either try to hit the top, or try to hit the bottom. Now, with that in hand, I think it becomes clear that there's a .5 chance that David *can* hit the top on the throw (since there's a .5 chance that he'll take a shot and hit the top), and likewise, for the bottom—in line with Success, and contrary to the control intuition.

So probability judgments defuse the Kenny/Boylan objection to Success. Unassertability doesn't entail falsehood. Once we control for the fact that ability ascriptions are generally weird when the agent could not try to do the action in question, probability judgments suggest that in pairs like Kenny's and Boylan's, one of the ability ascriptions in fact is true—we just don't know which.

It's worth noting that some people have reported that, while they agree about the judgments elicited in §4, their intuitions about probabilities in the Kenny/Boylan cases are unclear. I am not sure what to make of this. But even if judgments in the Kenny/Boylan cases are unclear for you, if you have the judgments elicited in §4, you should still agree that ability doesn't require control: if it did, Susie would have no chance at all of being able to hit the dartboard, and so on. Then there would remain a question about what makes the Kenny/Boylan cases different; but the answer is not 'lack of control', since if David lacks control over hitting the top of the dartboard, surely Susie lacks control over hitting a bullseye.

8 Non-agential ability ascriptions

In this final section, I'll argue that probability judgments also help resolve a different objection to a broadly conditional analysis.[13] The conditional analysis essentially involves the notion of *trying*: ⌜S is able to φ⌝ says that, if S *tries* to φ, she succeeds. The ACA still centrally invokes a notion of trying, as do all variants on the conditional analysis I know of.

However, there are cases where we apparently ascribe abilities to non-agents, as in (29) (from Irene Heim, attributed to Maria Bittner) or (30):

(29) This elevator is able to carry three thousand pounds.

(30) This black hole is able to absorb that galaxy.

I will argue here that probability judgments suggest that these cases are actually very different: (29) is an ability ascription, where the trying is done by a covert, generic agent, while (30) is a circumstantial modal. Neither is a problem for a conditional analysis.

[13] Thanks to Cian Dorr for suggesting this line of argument.

Start with (29). Suppose that I tell you that, conditional on loading the elevator with three thousand pounds of cargo, there is a 30% chance that the cord will snap, and a 70% chance that the elevator will work as normal. In that case, what's the probability of (29)? Intuitively, 70%. That is, credences in this case still seem to track conditional probabilities, in exactly the way that the conditional analysis suggests: the conditional probability of the elevator succeeding at carrying three thousand pounds, if you try to make it carry three thousand pounds. Of course, it's not the elevator that's trying. But (generic) *you* can try loading the elevator, and that seems to be what (29) is talking about: what happens *if you try*. That suggests an analysis of sentences like (29) along the lines of a conditional analysis, but with a covert generic agent.[14]

For another argument that we need an analysis along these lines, consider (31), from Melissa Fusco (p.c.):

(31) This knife can cut bread.

It just seems obvious that what (31) says is something about what a generic agent can do with this knife, not what the knife can do on its own. Whatever analysis you have of 'can' here, it plausibly will have to involve implicit agents of some kind.

Now turn to (30). Appealing to a covert generic agent obviously won't help here: the sentence clearly has nothing at all to do with agents, generic or otherwise, trying to do things. So this is, on the face of it, a harder case for any form of the conditional analysis. But now note that this case also seems totally unlike all the cases of ability ascriptions we've looked at so far vis-à-vis probabilities. In all the cases we've looked at, there is a very salient probability judgment about the ability ascription in question which matches a salient *conditional* probability judgment. But this doesn't seem to be true in this case. Suppose that the black hole has a 70% chance of swallowing the galaxy conditional on such-and-such physical processes taking place in the galaxy, and no chance otherwise. I don't see any way of filling in 'such-and-such' that makes it intuitive for your credence in (30) to be 70%.

What *should* your credence in (30) be? Well, it seems like it should just track your credence that there is *some* possibility that the black hole absorbs the galaxy. As always, there is context-sensitivity here, but (30) seems to just be saying that it is consistent with the black hole and galaxy's physical properties, and the laws of physics, that the former absorb the latter. Suppose for instance that you are sure that physical law and the black hole and galaxy's structure are consistent with the black hole absorbing the galaxy. Then it seems you should be sure of (30). Suppose instead that we are *unsure* what kind of black hole it is; your credence that it is big enough to absorb the galaxy is 70%. Then intuitively your credence in (30) should be 70%. Conditional probabilities don't seem to essentially enter the picture.

[14]This is something Mandelkern et al. (2017) suggest about cases like this.

Instead, the meaning of 'able' in (30) really seems to be that of an existential modal—the diamond of modal logic.

Given how polysemous modal words in general are (in English, as well as many other languages), it would be unsurprising to find that 'able' has readings where it is used as a circumstantial modal, in addition to those where it is used as an agentive modal. Adverting to polysemy like this would be theoretically unsatisfying if we were just using it to explain away counterexamples to a conditional analysis of the agentive reading. But probability judgments seem to provide clear evidence in favor of the hypothesis that there is something very different going on in (30) than in the cases we have looked at: these judgments suggest that, when 'able' is used to talk about scenarios where no agent is (or could be) involved, it is interpreted as an existential modal, along the lines of standard analyses of circumstantial modals; while when it is used to talk about agency, it is interpreted along the lines of the conditional analysis.

9 Conclusion

Many have thought that ability involves control, so that, say, for Susie to be able to hit a bullseye, hitting a bullseye must be somehow in her control. But probability judgments about ability ascriptions in cases like this show that this is wrong: ability is compatible with lack of control. They also provide a new argument for a conditional analysis of ability of some form: in many cases, judgments about the probability of ability ascriptions match judgments about the probability of the conditional which, according to the simplest form of the conditional analysis, corresponds to the ability ascription. Finally, these judgments provide a response to the two best arguments against any conditional analysis—first, the Kenny/Boylan argument against Success; second, the argument that we can ascribe abilities to non-agents. And more generally, probability judgments about chancy abilities provide a new source of evidence that any adequate theory of ability must explain.

Of course, there may still be indirect connections between ability and control. For one thing, being in a position to *assert* or *know* an ability ascription may often require knowing that the relevant agent has control over the relevant action. Likewise, *generic* ability ascriptions—the kind of thing we express with 'Susie is generally able to hit a bullseye', or 'Susie has the ability to hit a bullseye'—may involve control: indeed, to be *generally* such that you succeed if you try may just be what it is to have the action in question under control. But what probability judgments show is that these connections between ability and control, if they exist at all, are indirect: they are not encoded in the truth-conditions or presuppositions of ability ascriptions. Ability does not entail control; success does entail ability.

References

Adams, E. (1975). *The Logic of Conditionals*. Dordrecht.

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

Bacon, A. (2015). Stalnaker's thesis in context. *Review of Symbolic Logic*, 8(1):131–163.

Bhatt, R. (1999). Ability modals and their actuality entailments. In Shahin, K., Blake, S., and Kim, E.-S., editors, *The West Coast Conference on Formal Linguistics (WCCFL)*, volume 17, pages 74–87.

Boylan, D. (2020). Does success entail ability? *Noûs*.

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

Chisholm, R. M. (1964). J. L. Austin's philosophical papers. *Mind*, 73(289):1–26.

Cross, C. B. (1986). 'Can' and the logic of ability. *Philosophical Studies*, 50(1):53–64.

van Fraassen, B. (1976). Probabilities of conditionals. In Harper and Hooker, editors, *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, volume I, pages 261–308. D. Reidel Publishing Company, Dordrecht-Holland.

Fusco, M. (2020). Agential free choice. *Journal of Philosophical Logic*, 50:57–87.

Giannakidou, A. (2001). The meaning of free choice. *Linguistics and Philosophy*, 24:659–735.

Giannakidou, A. and Staraki, E. (2012). Ability, action, and causation: from pure ability to force. In Mari, A., Beyssade, C., and Prete, F. D., editors, *Genericity*. Oxford University Press.

Goldstein, S. and Santorio, P. (2021). Probability for epistemic modalities. *Philosophers' Imprint*, 33.

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

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

Holguín, B. and Lederman, H. (2022). Trying without fail. Manuscript, Johns Hopkins and UT Austin.

Horty, J. F. and Belnap, N. (1995). The deliberative stit: A study of action, omission, ability, and obligation. *Journal of Philosophical Logic*, 24(6):583–644.

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

Jaster, R. (2020). *Agents' Abilities*. De Gruyter.

Kaufmann, S. (2009). Conditionals right and left: Probabilities for the whole family. *Journal of Philosophical Logic*, 38:1–53.

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

Khoo, J. (2022). *The Meaning of* If. Oxford University Press.

Khoo, J. and Santorio, P. (2018). Lecture notes: Probability of conditionals in modal semantics. Unpublished manuscript.

Kratzer, A. (1977). What 'must' and 'can' must and can mean. *Linguistics and Philosophy*, 1(3):337–355.

Kratzer, A. (1981). The notional category of modality. In Eikmeyer, H. and Rieser, H., editors, *Words, Worlds, and Contexts: New Approaches in Word Semantics*, pages 38–74. de Gruyter.

Lehrer, K. (1976). 'Can' in theory and practice: A possible worlds analysis. In Brand, M. and Walton, D., editors, *Action Theory*, pages 241–270. D. Reidel.

Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.

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

Lewis, D. (1976b). Probabilities of conditionals and conditional probabilities. *The Philosophical Review*, 85(3):297–315.

Loets, A. and Zakkou, J. (2022). Agentive duality reconsidered. *Philosophical Studies*, 179:3771–3789.

Mandelkern, M., Schultheis, G., and Boylan, D. (2017). Agentive modals. *The Philosophical Review*, 126(3):301–343.

McGee, V. (1989). Conditional probabilities and compounds of conditionals. *The Philosophical Review*, 98(4):485–541.

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

Rothschild, D. (2013). Do indicative conditionals express propositions? *Noûs*, 47(1):49–68.

Santorio, P. (2022). Ability modals as causal modals. In Degano, M., Roberts, T., Sbardolini, G., and Schouwstra, M., editors, *Proceedings of the 23rd Amsterdam Colloquium*, pages 267–274.

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

Stalnaker, R. (1981). A defense of the conditional excluded middle. In Harper, W. L., Stalnaker, R., and Pearce, G., editors, *Ifs: Conditionals, Belief, Decision, Chance, and Time*, pages 87–104. Reidel.

Stalnaker, R. C. and Thomason, R. H. (1970). A semantic analysis of conditional logic. *Theoria*, 36(1):23–42.

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

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

Vranas, P. (2010). What time travelers may be able to do. *Philosophical Studies*, 150(1):115–121.

Willer, M. (2021). Two puzzles about ability *can*. *Linguistics and Philosophy*, 44:551–586.