PRINCIPIA **26**(2): 257–276 (2022) doi: 10.5007/1808-1711.2022.e80667
Published by NEL — Epistemology and Logic Research Group, Federal University of Santa Catarina (UFSC), Brazil.

ABILISM ASCENDANT?

KEVIN MEEKER
University of South Alabama, USA
kmeeker@southalabama.edu

Abstract. John Turri has recently called for a major shift in how the vast majority of philosophers think of knowledge. Instead of maintaining that knowledge must proceed from reliable processes, he urges epistemologists to move toward an "abilist" view that allows knowledge to proceed from abilities that are not truth-conducive. More strongly, he claims to have provided conclusive reasons for abandoning the idea that knowledge requires reliability. In this paper I explain why Turri has failed to make the case for preferring abilism.

Keywords: abilism • abduction • achievement • reliabilism • virtue epistemology

RECEIVED: 13/04/2021 REVISED: 23/02/2022 ACCEPTED: 16/03/2022

1. Ab initio

According to John Turri, contemporary epistemology needs a new beginning; for the discipline is hopelessly encumbered by its "central dogma" of "knowledge reliabilism" (Turri 2016, 190). While the reliabilist thesis that *epistemic justification* must be truth-conducive is controversial, Turri contends the vast majority of epistemologists simply assume that knowledge is truth-conducive. Challenging this deeply entrenched reliabililst view about knowledge, he argues that knowledge can be unreliable and proposes a successor view to guide our future investigations: "abilism". On this view knowledge is simply "true belief manifesting the agent's cognitive ability or powers" (Turri 2016, p.225).[1]

After offering this abilist proposal of knowledge as true belief arising from non-truth-conducive abilities, he provides a remarkably sweeping assessment of the state of the field:

> We have passed the point where it is reasonable to assume that reliabilism is true. Indeed, quite the opposite — moving forward we should assume that reliabilism is false. Potent theoretical and empirical considerations have converged to support the conclusion that unreliable knowledge is not only possible but actual and widely recognized as perfectly ordinary. A new day has dawned. It's time for the discipline to wake from its reliabilist slumber (Turri 2016, p.226).[2]

On the surface, this assessment seems rather harsh. Is it really *unreasonable* (passed the point where it is reasonable) to refuse to go along with this new abilist theory? Of course, Turri talks in terms of "assumptions"; and perhaps we should not make assumptions without considering arguments against such assumptions. But note that Turri appeals to "potent" considerations against reliabilism. How potent?

To answer this question, consider how Turri trumpets the demise of reliabilism after examining some results from experimental philosophy:

> When combined with the weakness of existing arguments for reliabilism and the recent emergence of well-supported alternative views [e.g., abilism] that predict the widespread existence of unreliable knowledge, the present findings [from experimental philosophy] are the final exhibit in a conclusive case for abandoning reliabilism in epistemology (Turri 2016, p.189).

If the case against reliabilism is *conclusive*, then presumably it would indeed be irrational for epistemologists to persist in their reliabilism. Given that the empirical findings to which he points are only the "final" exhibit in the conclusive case, though, the "theoretical" considerations presumably bear much of the weight of this debate.

What are these considerations? Turri provides the most comprehensive account of these arguments in his 2015a. In his own words:

> I have presented two arguments against truth-conducive reliabilism, one concerning achievements, the other concerning explanatory [abductive] inference. Each argument is valid with plausible and defensible premises. I submit that it is more likely that one of these arguments is sound than that truth-conducive reliability is true. It seems more likely than not that unreliable knowledge is genuinely possible. I conclude that the conventional wisdom is wrong about the relationship between knowledge and reliability (Turri 2015a, p.542–543).[3]

This present paper concentrates on these theoretical arguments from achievement and abduction. I contend that these arguments are more unconvincing than they are powerful. I leave aside Turri's empirical argument based on results from experimental philosophy because it raises a host of controversial issues about this new approach that would take us too far afield. If I can undermine his two theoretical arguments, then I will have provided strong evidence that Turri is incorrect to claim that we have *converging* considerations that constitute a conclusive case for abandoning reliabilism. Indeed, if reliabilism can emerge unscathed from such theoretical considerations, then it is doubtful that one will be able to build a convincing case for abilism on experimental philosophy alone — let alone a conclusive case.

The structure of this paper is as follows. In the next (second) section, I examine Turri's argument from achievements. The third section tackles further issues arising from the ways that Turri tries to support his argument from achievements. In the

fourth section, I discuss the argument from abduction. Finally, the fifth section not only sums up my contentions in defense of reliabilism and against abilism but also briefly sketches a reason to think that it is still acceptable to assume that reliabilism is true when theorizing about epistemology.

2. Achievements and Abilities

Building on recent epistemological discussions emphasizing that knowledge is a type of intellectual achievement, Turri argues that reflection on the nature of achievements reveals that knowledge does not require reliability (see, e.g., Turri 2012, p.186–187; Turri 2015b, p.320; Turri 2016, p.190–191; Turri 2018, p.313). By far the most developed version of this attack on reliabilism appears in the very beginning of the second section of his 2015a. There he explicitly states the argument in *modus ponens* form (Turri 2015a, p.531):

(1A) Achievements don't require reliable abilities.
(2A) If achievements don't require reliable abilities, then unreliable knowledge is possible.
(3A) Unreliable knowledge is possible.

In the course of the rest of that section, he elaborates on and even modifies his reasoning in several important respects. These modifications and elaborations ultimately sow doubts about the plausibility of this argument.

One problem that Turri explicitly confronts is that the argument as stated suffers from a scope ambiguity. One could interpret the argument very strongly as follows:

(1A*) No achievements require reliable abilities.
(2A*) If no achievements require reliable abilities, then unreliable knowledge is possible.
(3A*) Unreliable knowledge is possible.

On the other hand, we can interpret it more modestly:

(1A**) Some achievements don't require reliable abilities.
(2A**) If some achievements don't require reliable abilities, then unreliable knowledge is possible.
(3A**) Unreliable knowledge is possible.

Do either of these two represent Turri's intended argument?

Interpreted the first way, premise (1A*) seems questionable, as Turri admits. For (1A*) is a wide-ranging, exceptionless universal generalization that would be difficult to support with a simple argument.[4] On the other hand, if we opt for the second

argument, then we should doubt premise (2A**). As Turri puts it: "... the fact that some intellectual achievements can be unreliably produced doesn't lend much support to the claim that knowledge itself can be" (Turri 2015a, p.534). In response to the problems with these arguments, Turri proposes another version:

(1A***) Generally speaking, achievements don't require reliable abilities.

(2A***) If, generally speaking, achievements don't require reliable abilities, then unreliable knowledge is possible.

(3A***) Unreliable knowledge is possible.[5]

The phrase 'generally speaking' means that we are talking about "dominant tendencies, or what is typical, or what is natural and normal for a kind" (Turri 2015a, p.534).

I deny that (1A***) is plausible because Turri's assumptions about what is "normal" fail to take into account significant differences among different types of achievements and how reliable abilities might be related to these different types of achievement. More specifically, as we examine Turri's arguments and examples, I shall show how distinguishing among at least three types of achievements — adversarial achievements, alpha achievements and apex achievements — reveals some complexities in Turri's approach that otherwise might go unnoticed.[6] We should note that while these categories are not intended to exhaustively capture all types of achievements, they will aid us in understanding more broadly how Turri's examples and arguments are not nearly as convincing as he claims that they are.

Turri supports (1A***) with several examples. One of the most prominent is that of Geno, an eleven month old who "... has been daily gaining confidence in his ability to balance on two feet. Today he stood holding onto the couch when an object in Dad's hand commanded his attention. Inspired, Geno let go, and... took his first two steps" (Turri 2015a, p.531). From this scenario, familiar to many parents, Turri draws a momentous conclusion. Once again, in Turri's words: "Those first steps were an achievement... But of course at eleven months old he's still a highly unreliable walker. So achievement can issue from even highly *un*reliable ability" (Turri 2015a, p.531). Turri generalizes that "Achievements populate the road to proficiency in many spheres" (Turri 2015a, p.532; see also Turri 2018, p.313). These examples are trickier than Turri admits. For I doubt that such achievements issue from highly unreliable abilities — at least not in any sense that is helpful to Turri's defense of (1A**).

In line with my distinctions about different types of achievements, let us call these alpha achievements. Roughly speaking, an alpha achievement is, as the name suggests, a "first" achievement. Walking is normal and usually not considered an achievement because, presumably, it is so "common" or mundane. But children's first steps are achievements, as Turri claims. Does it follow, though, that alpha achievements issue from highly unreliable ability?

If we understand reliability in terms of track record, then perhaps Turri has a

strong case. The child has a (short) lifetime of being unable to walk. Measured against this time period, a few steps is not enough to claim that child is reliable at walking. But a track record understanding of reliability in this case seems to miss the mark. For very quickly, children become reliable walkers. If we measured reliability by track record over the appropriate lifetimes, then we have to deny that these children now walking very consistently can reliably walk. But clearly they can. If we understand reliability in a propensity manner, then the children have the ability to reliably walk.

It seems, then, that Turri appropriately assumes that, at first, the claim that children cannot reliably walk means that the children "could not reliably produce such results" (Turri 2015a, p.532.) But if a child cannot reliably reproduce these movements, should we really claim that the child has taken "steps"? After all, we can distinguish between walking and stumbling forward. When children take what seem to be their first steps, are they walking/stepping or merely stumbling forward in a way that approximates walking? How could one tell the difference? If children can consistently stand and take a few steps right after what appear to be their first steps, then I agree that they are walking but would also insist that their first *steps* also revealed that they had the ability to *reliably* walk. So their alpha achievement did *not* issue from unreliable abilities — let alone highly unreliable ones.

Turri could respond that because they can only consistently take a few steps, they have not yet gained the ability to walk reliably. But such a response would confuse the ability to walk reliably with the ability to walk long distances reliably. Just because some lack the stamina to walk long distances, it does not follow that they cannot reliably walk.

Turri seems to assume that first steps are not consistently followed by other steps. So the walkers are not reliable. To return to our former distinction, in such a case it seems more plausible to claim that such movements were not really steps at all, but merely instances of stumbling that approximate walking — not walking itself. More strongly, it is plausible to think that one feature that distinguishes walking/stepping from stumbling forward is some kind of reliable ability. If one does not possess a reliable ability to walk, then simply putting one foot in front of another is only "walking" by accident. It is not really walking.

To explore this idea further, consider some analogies. Suppose you are teaching Deanna to dive even though she has never even jumped off of a diving board. Suppose further that the first time she jumps off the board and happens to enter the water headfirst. The next one hundred times, she goes in feet first. The first attempt was not really a dive, it was simply a jump that happened to approximate diving. Granted, the first jump can still be an achievement; but the achievement is showing the courage to jump. Having overcome her fear of hitting the water, Deanna can continue to try to dive. Similarly, Geno's first stumbles could be an achievement that breaks him of his fear of letting go and falling so that he can practice until he takes his first steps,

instead of merely stumbling forward. But then how precisely should we describe the very first dive or the very first step?

To answer this question, consider another analogy. When first learning to ride a bike, the key is balance. One can "wobble forward" by frantically turning the front wheel from side to side for a short period. But at some point, how to balance "clicks" and one is able to ride the bike instead of wobbling forward. This "clicking" is an instantaneous — or near instantaneous — acquisition of the ability to ride a bike reliably. The same is true, I suggest, of walking (or diving). Part of what makes an alpha achievement an achievement, I suggest, is that the person is exhibiting for the first time a newly acquired reliable ability.

I suspect that Turri would disagree. I also suspect that sorting through this disagreement would lead to difficult discussions about how quickly stumbling forward can morph into walking or how quickly wobbling forward can morph into riding. Without knowing exactly how Turri would respond to such distinctions, it is difficult to know how to argue the point much further.

While some may not agree with my assessment with these alpha achievements, I hope to have provided reason to think that Turri's quick use of them is not as convincing as he thinks it is. And even if he is right that alpha achievements do not require reliable abilities (which I have denied), these alpha achievements seem to be a small percentage of achievements; what's more, there is no reason to think that we should consider knowledge to be a type of *alpha* achievement. To put this point another way, for this example to support Turri's attack on the idea that knowledge is a form of reliable achievement, we would need some reason to think that this notion of an alpha achievement relevantly overlaps with the reliabilist concept of achievement with respect to knowledge.[7] So let us look at some of Turri's other examples of achievements to see if they fare any better.

3. Apex and Adversarial Achievements

Perhaps the most developed example that Turri uses, and one to which he frequently returns, is that of the famous baseball player Ted Williams (see Turri 2015a, p.531, 533, 535 and 541; Derek Jeter is briefly mentioned as a similar example in his 2012). Let us examine exactly how Turri puts the argument: "Ted Williams is the best baseball hitter ever. But he normally fails to get a hit. The relevant ability could at best be counted on to produce a hit about four in ten times: his best yearly average was .407 and his lifetime average much lower" (Turri 2015a, p.531). Before we examine how the notions of apex and adversarial achievements help us to examine this type of example more thoroughly, we first need to question some basic assumptions of Turri's assessment.

Turri apparently assumes that the relevant measurement of batting ability is batting average. But while batting average, which is determined by dividing the number of hits by the number of at bats (not simply plate appearances), is one of the most well-known statistics, it is misleading to focus on it in this discussion. For many baseball aficionados consider on-base percentage a better, albeit imperfect, measure of the ability of a batter. This statistic counts not only hits but also the number of times a batter reaches base by walks, or being hit by a pitch. (Reaching base because of a fielding or throwing error actually hurts one's on-base and batting average.) In the year to which Turri refers Ted Williams had an on-base percentage above 55%. So Williams did not normally fail to get on base. He normally reached base.

Another important measurement of batting ability is slugging percentage. The formula for this statistic is as follows: ([singles] + [doubles x 2] + [triples x 3] + [home runs x 4])/[at bats]). This measurement recognizes that some hits are more of an achievement than others. Home runs are more impressive than, say, a single. So looking at players' batting averages for a game will not differentiate a player hitting four home runs in a game from one who hits four singles even though the former should clearly be credited with a bigger achievement. To put this point another way, if a position player hits the ball only one out of five at bats — and these hits are all singles — then, unless that player is a once-in-a-lifetime fielder, he almost certainly will not have a long career. But if all those hits are home runs, then that person would easily be considered the best player of all time even with such a low batting average. One reason that Ted Williams is considered by many to be the best baseball hitter ever is because his lifetime slugging percentage was above 63% (second only to Babe Ruth, albeit by a fairly significant margin).

But let us ignore some of the complications that surround slugging percentages and trying to measure the value of various types of achievements and return to on-base percentage.[8] This statistic is an imperfect measurement of one's ability because one can *succeed* as a batter (perhaps better: achieve a successful at-bat) even if one does not get on base. A sacrifice fly, for example, does not help one's batting average (and actually hurts one's on-base percentage) despite the fact that it is universally considered a successful plate appearance. Similarly, to move a runner from second to third and thus be in a better scoring position (especially if there are no outs), a batter can hit the ball on the ground to right side of the infield and be retired. This "moving the runner" is again seen as a *successful* at bat even though it hurts one's batting average, slugging percentage and on-base percentage. By the same token, consider a batter who gets a full count (three balls and two strikes) and fouls off ten consecutive pitches before striking out. A sixteen pitch at bat is viewed as an extraordinarily successful at bat even if it damages ones's batting average and on-base percentage because it wears out the pitcher. In general, sacrifice flies, sacrifice bunts, moving a runner, and going deep in a count are considered quality plate appearances. Unfor-

tunately, because there is no agreed upon definition of when, precisely, a batter, say, goes "deep in the count", there is no official quality at bat statistic. But this does not stop teams from developing their own somewhat subjective, unofficial measurement. Even granted the vagueness of such a measurement, it stands to reason that many players are reliable in producing quality at bats.

Turri could object that without an official statistic, it is unclear how many baseball players actually are reliable in producing quality at bats. Moreover, although he constantly refers to Ted Williams, he points out that " ... hits by many lesser hitters are no doubt achievements too..." (Turri 2015a, p.531). So he could argue that lesser hitters produce achievements even if they do not reliably produce quality at bats.

These responses would be insufficient. When Turri simply claims that even the best hitter of all time normally failed to get a hit in even his best season, that seems to support (1A***). But when we realize that the goal of a hitter is not simply to get a hit, this simple point becomes questionable at best. For we have seen that hits do not exhaust the achievements of a batter, and a batter can achieve (or succeed) without getting a hit.

We could continue indefinitely debating whether Turri's notion of how to measure a successful plate appearance adequately captures the goal of a baseball batter. At this point, though, we have discussed enough issues that we are now in a better position to see how the notions of apex and adversarial achievements can shed further light on our dialectical situation. Let us start with the notion of an apex achievement. Crucial to Turri's example is Ted Williams' status not as any baseball player but a Hall of Fame major leaguer. His batting average is what I would call an apex achievement. An apex achievement is an achievement that somehow separates one in a positive way from others who can achieve similar things. Having the highest batting average in the major leagues — as opposed to say a local Little League — is an apex achievement. Of course, this type of achievement differs significantly from an alpha achievement. For, generally speaking, alpha achievements are what we might call synchronic achievements while apex achievements are more likely to be what we might call diachronic achievements. That is, alpha achievements such as a first step or first official hit are instances of a larger set of actions or achievements that as a whole may distinguish one person from another. One of the highest career batting averages would be a diachronic achievement. And because knowledge is arguably a synchronic achievement, one might wonder about the relevance of discussing these diachronic examples in the context of a discussion of knowledge.[9]

In response, we should note two points. First, in asserting (1A***), Turri seems to be making a claim about all types of achievements, both synchronic and diachronic. Second, even if he did mean to restrict his argument to synchronic achievements, many apex achievements seem to be synchronic achievements. Being able to play Bach concertos almost flawlessly every time, as some musicians can, is an apex achieve-

ment. So presumably is any particular instance of such a musician playing a flawless Bach concerto. In short, apex achievements include synchronic and diachronic instances.

One difference between the musician and the baseball player is that, generally speaking, baseball players' achievements occur in an adversarial environment. Adversarial achievements, as the name suggests, occur when someone (or something) is trying to prevent you from succeeding. A musician's synchronic apex achievement of another flawless performance of a Bach concerto is not an adversarial achievement — even if it is in the midst of a musical competition. For no one is doing anything actively to prevent a musician from succeeding; the judges are merely trying to discern whose success is the best.

The same can most assuredly not be said of baseball players. Note that hitting a moving baseball is an ability, and an impressive one at that. If we set up a pitching machine on a mound and programmed it to throw 85 mph, then very few people in the world could consistently hit the ball in fair territory. Major league players could do that over 90% of the time. They represent the top .01% of people in the world in terms of this ability.[10] They have reached the level of apex achievements. People at this level clearly have the ability to hit a baseball reliably, just as some musicians can play Bach concertos almost flawlessly every time. In other words, each particular hit of a player at this level results from reliable abilities. Of course, when playing against others in that top group, with the pitchers able not only to throw well over 90 mph as well as deceive you with slower breaking balls but also able to rely on eight other players in the field trying to record an out, then it becomes very difficult to record an *official* "hit". But an official hit is an adversarial achievement, earned in a difficult environment that stretches the ability of all players involved. Moreover, in an important sense every hit in a major league baseball game is an apex achievement because so few make it to this level.[11] Achievements in such circumstances do not show that the major league players lack the ability to reliably hit the baseball. One needs excellent bat speed, hand-eye coordination, and so on. Because (1A***) entails that, generally speaking, adversarial baseball achievements in the major leagues do not require reliable abilities, (1A***) is extremely implausible.

These examples also show that (1A*) is even more implausible because it is such a strong modal claim. Recall, though, that while Turri does not think (originally at least) that he needs it for his argument, he still endorses it because "... anything that can be reliably achieved can be unreliably achieved" (Turri 2015a, p.534). Interestingly, when referring to this article in his *Stanford Encyclopedia of Philosophy* entry, Turri summarizes the achievement argument this way: "... if knowledge is an achievement, then we should expect it to not require reliability, because no other achievement requires reliability" (Turri, Alfano, Greco, 2018). *No* achievement requires reliability?

This stance is exceedingly difficult to maintain. For even Turri admits that one needs well over a 90% fielding average to be considered a good fielder in baseball (Turri 2015a, p.540). Almost all major league players have such a high fielding average. Such apex achievements in an adversarial environment clearly require reliable abilities on *any* way of interpreting reliable abilities. Even if there are distant possible worlds in which one accidentally (i.e., without reliable abilities) fields a ball cleanly 99% of the time, in those worlds one is not *fielding* a baseball expertly — the ball is just (strangely and *accidentally*) bouncing into the correct place. Because I categorically deny that one can unreliably achieve a 99% fielding percentage, I also deny the sweeping claim that *anything* that can be reliably achieved can be unreliably achieved. Put another way, contrary to a clear implication of Turri's assertion, it is not possible for someone to unreliably achieve a 99% fielding percentage. Thus, we should all agree that some achievements require reliability. And I find nothing in Turri's argument that touches such a point.

Because Turri initially admits that he does not need (1A*) for his argument, perhaps an abilist could fall back on the claim that achievements *typically* don't require reliable abilities even if *some* achievements (playing a Bach concerto almost flawlessly; scoring a 99% fielding percentage over the course of a season) do. But I have argued that his examples fail to make the case that any achievements can issue from unreliable abilities. If I am correct that Turri has not produced a single clear case of an achievement that is unreliably produced, then his claim that achievements typically issue from unreliable abilities is in serious jeopardy.

To further defend reliabilist interpretations of Turri's examples, let us examine a similar example focusing on the real issue: knowledge. Suppose that a group of cognizers are competing to see who can spot Waldo most effectively in a series of pictures. Hundreds of competitors are shown pictures and given 10 seconds to spot Waldo. Almost all of the competitors spot Waldo 10-20% of the time. But Amanda wins the competition because she and she alone spots Waldo 40% of the time. Amanda is the Ted Williams of the *Where's Waldo?* world. Clearly, her spotting Waldo 40% of the time is an achievement.

But what about the individual spottings of Waldo — are these knowledge?[12] That is, when Amanda identifies Waldo does she *know* that it is Waldo that she has seen? While she may not be able to articulate how she finds Waldo so much more quickly than other people, she clearly knows Waldo when she sees him. If anyone has Waldo knowledge, it's Amanda. Now Turri would point out that Amanda fails to identify Waldo more often than she identifies him and suggest that these failures mean that Amanda's knowledge is unreliable. As Turri puts it, the reliabilist thesis is that "… knowledge must be produced by *a reliable ability that yields mostly true beliefs*" (Turri 2015a, p.535; Turri's emphasis). But how, exactly, should reliabilists relate "reliability" with "ability that yields mostly true beliefs"?

Although Amanda may not be able to achieve the knowledge reliably, it does not follow that the knowledge results from unreliable abilities — at least not in any sense that would undermine reliabilism. For it's not that she erroneously confuses Waldo with other drawn characters and thus when presented with a character can only tell whether it is Waldo or a non-Waldo less than half the time. Given the adversarial conditions (the time limit), she simply does not always have the time to identify Waldo. Her beliefs about the location of Waldo are all true; in the cases where she "fails", she does not have false beliefs — she simply lacks any belief about Waldo's location. Clearly, Amanda has an ability that yields mostly true beliefs in the sense relevant for knowledge. Overall, then, reliabilists can quite plausibly claim that they are not requiring that a cognizer achieve knowledge reliably — only that the knowledge issues from reliable abilities in the relevant sense. In short, then, none of Turri's achievement examples should lead us to expect that knowledge issues from unreliable abilities in the relevant sense.

4. Arguing about Abduction

In addition to this achievement argument, Turri also offers another theoretical argument, which suffers from a similar problem. The basic idea of the argument is this: abductive reasoning in science yields knowledge even though many such inferences are not reliable. He starts off his discussion with a scientific example: "We now know that space-time is curved. We know this because the space-time hypothesis best explains a wealth of data" (Turri 2015a, p.536). In the next four paragraphs (and a pie chart), he launches into a brief description of abductive reasoning, or what is also called inference to the best explanation, to show that unreliable knowledge is possible. Unlike the achievement argument, though, Turri presumably only needs one case of an unreliable abduction yielding knowledge to make his case. That is, Turri previously argued that because achievements typically don't require reliability, we should not expect the achievement of knowledge to require reliability either. In this abduction argument, Turri assumes that he needs to provide just one example of unreliable abduction *actually* leading to knowledge to show that unreliable knowledge is *possible*.

Strangely, his four sentence "discussion" of relativity theory does not explicitly claim, as far as I can tell, that the abductive inferences supporting relativity theory are unreliable. But then Turri simply states: "Inference to the best explanation yields knowledge if the explanation we arrive at is true. But even when it is true, the best explanation might not be very likely. So our disposition to infer to the best explanation might not be very likely" (Turri 2015a, p.536). Moving to an even more abstract level, Turri considers an unnamed hypothesis (H) that is forty-nine percent likely

and explains certain data (D) better than 51 other competing hypotheses that are each one percent likely. Turri then asserts: "H is by far the best explanation... even though it's more likely that one of ∼H1 through ∼H51 explains D... In such a case, it's reasonable to accept that H explains D. And if it's true that H explains D, it seems that you can thereby know that H explains D" (Turri 2015a, p.537). Turri's brisk survey of relativity theory and abductive reasoning leaves many questions unanswered.

Is what's true in this abstract case true of all abductive inferences? That is, do all abductive inferences that arrive at truth qualify as knowledge? Do true abductive inferences "typically" yield knowledge — just as achievements (allegedly) typically issue from unreliable abilities? If you know that H explains D, does it follow that you know that H? Is his description of how relativity theory is confirmed accurate? Did abduction *justify* relativity theory or was it merely operative in the *discovery/formulation* of the theory? Perhaps most basically, is his description of abduction as a legitimate inference in science accurate? Pausing briefly on this last question, we should note that philosophers of science such as Bas van Fraassen (see, for example, his 1989) have argued that abduction is an illegitimate form of scientific inference. Needless to say, those who think that van Fraassen has a good case — even if they do not ultimately agree with him — will wonder how Turri can build a convincing case without addressing such a fundamental and well-known debate. To put this point another way, while Turri does an exhaustive job of citing a wide range of epistemological literature supporting the idea that there is a consensus about knowledge requiring reliability, he ignores an important controversy related to abduction. Without addressing such a controversy, Turri's argument and example are incomplete at best.

Perhaps Turri means to avoid such questions and controversies, when, in the next eleven paragraphs, he focuses on a "case-study to help make the point vivid" (Turri 2015a, p.537). Interestingly, he explicitly states his argument in the context of a case-study not of actual scientific practice, but of a fictional TV show called *House*. On this show a doctor (Gregory House) attempts to diagnose patients with very rare illnesses. Reasoning abductively, House usually arrives at the correct diagnosis after several missed diagnoses. Turri claims that when House arrives at the correct diagnosis, he acquires knowledge (presumably before the diagnosis is confirmed). So Turri's explicit statement of the argument once again takes on a *modus ponens* form (Turri 2015a, p.539):

(1B) If House knows, then unreliable knowledge is possible.
(2B) House knows.
(3B) So unreliable knowledge is possible.

While Turri recognizes a scope ambiguity in his achievement argument, he overlooks a different type of ambiguity that plagues his choice to focus on a fictional (series of) abduction(s).

Consider a parallel of the first premise:

(4) If Marty McFly meets his parents years before he was born, then time travel is possible.

We can interpret this premise in at least two ways:

(4A) If Marty McFly meets his parents years before he was born *in the* Back to the Future *story*, then time travel is possible *in the* Back to the Future *story*.
(4B) If Marty McFly meets his parents years before he was born *in the* Back to the Future *story*, then time travel is possible in *the actual world*.

(4A) seems true, but only in a trivial sense that would not support any significant conclusion. (4B), on the other hand, allows us to construct an argument with a rather momentous conclusion. (For those unfamiliar with the movie: the antecedent is true.) But this argument with momentous import also raises immense controversy. Being able to tell a story about a time traveling DeLorian with a flux capacitor does not show that time travel is possible in the actual world. By the same token, we can interpret (1B) in at least two ways:

(1B*) If House knows *in the TV drama*, then unreliable knowledge is possible *in the TV drama*.
(1B**) If House knows *in the TV drama*, then unreliable knowledge is possible *in the actual world*.

Clearly, Turri needs (1B**) for his argument to work because the conclusion certainly seems, more exactly, to be (3B*) Unreliable knowledge is possible *in the actual word*.

So an initial question we must face is the following: why should we look to a fictional example to guide our judgments about possibility? We should note that philosophical thought experiments are often criticized because they are bad stories or underdescribed. The advantage of appealing to such widely known stories as popular TV shows is that they present many extensive, vivid details that have more potential to repay careful reflection. Although I think that such a development is healthy for philosophy in general, in this case I am unconvinced by Turri's argument. Because Turri never recognizes that he needs (1B**), he never explicitly addresses this problem.

The closest Turri comes to broaching this issue is the following comment in support of (2B), which he claims is: "... supported by intuition, and by the fact that millions of viewers, including trained epistemologists, detect no incoherence in the story line, week after week, over many seasons" (Turri 2015a, p.539). Unlike, say, the *Back to the Future* story, Turri could claim that an extended story in which House knows is "coherent" whereas the time travel scenario is "incoherent". Whether or not

one thinks the narrative is coherent will presumably depend on whether one thinks that House knows. In other words, (2B) or not (2B) is the real question.

Unlike Turri, though, I have no intuition that House knows. Granted, I have only watched a few scattered scenes of the show over the years and do not have any in-depth knowledge of the series. Nevertheless, my judgment is supported by Turri's own description of the show. According to Turri, in one episode House himself seems to deny that he has knowledge: "I'm *almost* always *eventually* right. You have no way of knowing when *eventually* is" (Turri 2015a, p.538; Turri's emphasis). House's assessment seems correct. To see why, consider again Amanda. Suppose that we had seven pictures to show her — each with a well-hidden Waldo. If she has ten seconds to look at each picture, given her 40% ability described above, it is highly likely (over a 97% chance) that she will *eventually* diagnose a Waldo in one of the pictures shown to her. But before she examines a particular picture, we do not know if she will correctly diagnose the whereabouts of the Waldo. By the same token, although it is highly likely that House will eventually reach the correct diagnosis, we do not know which diagnosis is going to be the correct one. Nor does House. It seems to me that he clearly lacks knowledge.

At one point, Turri seems to anticipate an objection along these lines. He considers the idea that "... the relevant method isn't *inference to the best explanation*, but rather *trying to solve the case*" (Turri 2015a, p.539). Turri resists this objection because he identifies trying to solve the case with the method of inference to the best explanation (Turri 2015a, p.540). Here I think Turri is mistaken in several respects.

First, because of the structure of the TV drama format, the audience knows that House's first diagnosis will be wrong and the one close to the end of the show will be right. House knows how to solve the case before one dies; but I would insist that the characters have no knowledge which diagnosis will be correct — unlike the viewing audience because of the structure of the show. In other words, his reliably solving the case before a patient dies, when coupled with the time structure of the show, provide the audience with an artificial "tell" that is absent in the actual world. House finally gets it right *just before the patient dies* — what are the odds?[13]

Second, to the extent that House knows, he seems to be on firmer — or more reliable — ground than mere unreliable inference to the best explanation. Consider Turri's description of another episode "Occam's Razor":

> House's team is left to decide between two potential diagnoses, described as "a ten-million-to-one shot" versus " a million-to-one shot": House settles for the latter. Would even *this* diagnosis yield knowledge if true? It can easily seem like a stretch to say 'yes', and one suspects that the show's writers are here deliberately exaggerating one of the show's central themes, namely, intellectually overcoming unfavorable odds to save a patient's life (Turri 2015a, p.539).

Strangely, I have the opposite reaction to this situation based on Turri's description of it.

If House truly does have only two diagnoses from which to choose, then presumably all other diagnoses have been eliminated (with, I suppose, the exception of a previously unknown medical condition). In such a case, while it may be that only one in one million people who present such symptoms have this condition, when all other conditions have been eliminated except the one in ten million condition, obviously the former diagnosis is very likely to be correct. If some object that there are other possible diagnoses, then it is at best unclear to me why the situation would be described as a choice between *two* diagnoses.

If I understand the situation, then, the abductive inferences become more reliable as an episode progresses because the number of competing exotic hypotheses dwindles and information about the state of the patient increases. In other words, perhaps House's method is not just abduction but also a process of elimination or trial and error which could eventually lead to a diagnosis that has a high likelihood of being correct. So even if House has knowledge in some episodes, it is reliably produced.

To reiterate, though, my inclination is simply to deny that House knows. Granted, my lack of familiarity with the show leads me to admit that perhaps House knows because he follows a process of elimination that ends up being reliable. For what it's worth, though, I have discussed this issue with dozens of viewers of House and they — philosophers and nonphilosophers alike — have *unanimously* rejected (2B). Based on these conversations, I am quite puzzled about how Turri can so confidently assert that "millions" of viewers find no incoherence in saying that House knows. Be that as it may, I have also argued more substantively that even Turri's description of the show, which is explicitly presented to undermine reliabilism, fails to be convincing. For his description reveals either that House does not know or that House at least sometimes uses a reliable method that includes something like a process of elimination.

Suppose that the show — or the writers of the show — really do mean to assert that House knows even though he is using an unreliable abductive method. I would then contend the show exhibited a type of aesthethic or narrative incoherence. In other words, I would accept that House knows in the TV show but reject (1B**), the undefended premise Turri needs for his argument to work. To put the point another way, just as I don't think that watching the *Back to the Future* movies would convince those skeptical of time travel that time travel is possible, so watching *House* would not convince me that unreliable knowledge is possible.

Although I deny that House has knowledge, it might be helpful to briefly consider the House case in terms of his achievements, the notion that played such a large role in Turri's previous argument. House's ability to save patients' lives is, let us grant, an achievement. Even if 80% of his patients died, I would still grant that House's diagnoses are achievements. Why? Recall the notion of a batter's slugging percentage

discussed in a previous section. If you know your baseball (or if you did the math), you would know that it is possible for someone to have a slugging percentage that exceeds 100%. Such a mathematical possibility might make it seem as though it is not a good way to think of Turri's point because you can't have greater than 100% of truth when it comes to beliefs. Nevertheless, there is a sense in which some truths are more important than others just as some hits are more important than other hits. House's ability to "diagnose the undiagnosable" — even if it only saves 20% of his patients — is more important and more impressive than a doctor who can diagnose, say, a cold with 99% accuracy. In other words, House's "slugging percentage" is incredibly high for a medical doctor even if his "batting average" is lower than those who deal with more run of the mill cases. And we can agree that this apex diagnostic achievement does not mean that House lacks relevant reliable abilities any more than a home run hitter with an apex slugging achievement lacks relevant reliable hitting abilities because of a "low" batting average. But neither House nor the home run hitter knows that a diagnosis will work or that an at bat will produce a home run until one makes contact with the ball or sees the result of the diagnosis.

Before closing this section, we should note that Turri's argument from abduction has several positive features. First, as mentioned previously, focusing on an actual story — such as a TV show — has the potential to provide thought experiments that are more detailed and probative than the typical thought experiments philosophers devise. Second, because epistemologists and philosophers of science have historically not had much interaction, any attempt to bring the two together is welcome. Unfortunately, in this case the argument fails to capitalize on these advantages. Turri's brief reference to science fails to engage (any of) the relevant debates in philosophy of science. Moreover, his own description of the House story fails to give a coherent picture of how House knows using unreliable abduction. Regarding his explicit argument, then, I deny (2B) — and any temptation I have to accept (2B) would cause me to reject (1B**), which is clearly needed for a valid argument.

5. Assuming anti-Abilism?

While I have agreed that Turri's two theoretical arguments against reliabilism are valid on some suitable interpretations, I have contended that they are not sound. More specifically, I have argued that when we consider the features of different types of achievements — alpha, apex, and adversarial — we can see that Turri has failed to give a convincing example of an achievement that issues from unreliable abilities. Even if one questions this contention, it seems extraordinarily implausible to claim that achievements *typically* issue from unreliable abilities. Even more implausible is the suggestion that no achievements require reliable abilities. With regard to his

argument from abduction, I have contended that his discussion fails in the case of the fictional medical example and is incomplete at best with regard to issues surrounding actual scientific practice. Both arguments are defensible in the sense that Turri has defended them and I have not argued that it is impossible that Turri will be able to respond to the criticisms I have lodged. Nevertheless, given the problems I have pointed out with his "potent" theoretical arguments, I deny that it is more likely that one of these arguments is sound than that reliabilism is true.

Without strong support from these theoretical arguments, I doubt that the empirical argument is going to be able to support a conclusive case against reliabilism. Likewise, given the controversy surrounding experimental philosophy in general, it is difficult to see how Turri can maintain that it is unreasonable to continue to be a reliabilist. Of course, Turri has argued that "Philosophers have simply assumed that knowledge requires reliability, without offering any evidence in support of this assumption" (Turri 2015b, p.320). Could it be that evidence-free epistemological assumptions are trumped by Turri's empirical argument so that the balance of evidence does indeed favor abilism?

To be sure, we cannot rule this possibility out *a priori*. But given what we have learned in our discussion of Turri's theoretical arguments, I want to suggest how moving forward we can be quite reasonable in assuming that reliabilism is true. To elaborate, perhaps there is a sense in which some philosophers have "simply assumed" reliabilism; but consider in this context the following point from Linda Zagzebski: "The generality problem is a problem for any theory that makes reliability a necessary condition for knowledge, including my theory... so we should look now at a way to handle the problem" (Zagzebski 1996, p.309). Clearly Zagzebski is not simply assuming reliabilism as obviously true if she thinks that such theories confront difficulties, such as the generality problem. In other words, seeing the need for handling problems with a theory shows that one is not just ignoring the evidential status of a theory. Indeed, to the extent that a theory successfully handles such problems, we can plausibly be said to have evidence that it is true.

Take Turri's theoretical concerns about abduction and achievements. By showing how reliabilism can handle such "potent" problems, I have shown the durability and coherence of reliabilism. Insofar as durability and coherence are considered theoretical virtues, and such virtues are reasonably said to be epistemic virtues, we can say that we have reason to accept reliabilism. Unfortunately, we lack the time to pursue such matters in depth in the present context; my main point, though, is that reliabilism's ability to withstand Turri's attack could be plausibly viewed as a reason to accept reliabilism. So until Turri develops a better theoretical case against reliabilism and for abilism, I suggest that it is appropriate in at least some contexts for epistemologists to assume that reliabilism is true. As we have seen, Turri himself is not opposed to assumptions in future epistemological theorizing; he just thinks that it is

abilism that should be assumed moving forward. I disagree.

Does this mean that Turri is unreasonable to pursue his abilist research or even believe that abilism is true? Not at all. But his theoretical arguments for abilism and against reliabilism are not nearly sufficient to conclusively show that contemporary epistemologists as a whole should abandon reliabilism in favor of his abilist hypothesis.

References

Bradford, G. 2015. Knowledge, Achievement, and Manifestation. *Erkenntnis* **80**(1): 97–116.

Greco, J. 2020. Knowledge-producing Abilities. In: C. Kelp and J. Greco (ed.), *Virtue-theoretic Epistemology: New Methods and Approaches*, p.124–146. Cambridge: Cambridge University Press.

McKinnon, R. 2014. You Make Your Own Luck. *Metaphilosophy* **45**(4-5): 558–577.

Schoenfield, D. 2019. Generation OPS: How powerful young All-Stars are redefining the game. www.espn.com/mlb/story/_/id/27148938/generation-ops-how-powerful-young-all-stars-redefining-game. Access: 22.03.2022.

Turri, J. 2012. Review of *Achieving Knowledge: A Virtue-Theoretic Account of Normativity*, by John Greco. *Mind* **121**(48): 183–187.

Turri, J. 2015a. Unreliable Knowledge. *Philosophy and Phenomenological Research* **90**(3): 529–545.

Turri, J. 2015b. From Virtue Theory to Abilism: Theoretical and Empirical Developments. In: C. Miller, R. M. Furr, A. Knobel and W. Fleeson (ed.), *Character: New Dimensions in Philosophy, Psychology, and Theology*, p.315–350. Oxford: Oxford University Press.

Turri, J. 2016. A New Paradigm for Epistemology: From Reliabilism to Abilism. *Ergo* **3**(8): 189–231.

Turri, J. 2018. Virtue Epistemology and Abilism on Knowledge. In: H. Battaly (ed.) *The Routledge Handbook of Epistemology*, p.309–316. New York: Routledge.

Turri, J.; Alfano, M.; Greco, J. 2018. Virtue Epistemology. In: E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*. Summer 2018 Edition. https://plato.stanford.edu/archives/sum2018/entries/epistemology-virtue/. Access: 22.03.2022.

Van Fraassen, B. 1989. *Laws and Symmetry*. Oxford: Oxford University Press.

Zagzebski, L. 1996. *Virtues of the Mind*. New York: Cambridge University Press.

Notes

[1]Perhaps more accurately, he defines knowledge as "approximately true thin belief manifesting cognitive ability" (Turri 2015b, p.321). For some criticisms of Turri's appeal to "manifesting", see Bradford 2015. In his 2015a, Turri labels his view "ecumenical reliabilism" and cashes out reliability in terms of either truth-conducive reliability or non-truth-conducive trustworthiness of the processes that produce beliefs. His "abilism" view seems to be more developed in terms of what it means for a process to be "trustworthy".

[2]More strongly than contending that reliabilism is false, he also confidently claims that his abilist view "captures the ordinary concept of knowledge and is the clearest, most accurate, and informative philosophical definition available" (Turri 2015b, p.326). To be clear, if Turri is correct that knowledge does not require reliability, then epistemic justification cannot require reliability either insofar as it is necessary for knowledge. My defense of reliabilism in this paper is only meant to support the claim that knowledge, not epistemic justification, requires reliability. Consequently, when I refer to "reliabilism" throughout the paper, I mean "knowledge reliabilism". More precisely, I am defending the view that belief can count as propositional knowledge only if it arises from some kind of reliable ability/process because I believe that this is the view that Turri is attacking. [My thanks to an anonymous referee for comments that prompted this clarification.]

[3]Elsewhere, Turri similarly claims that the "best" argument for knowledge reliabilism is "equivocal" and refers to his 2015a for "positive arguments for the possibility of unreliable knowledge" (Turri 2016, p.190). After briefly rehearsing the achievement argument, he claims that "At this point, the balance of evidence has tipped against reliabilism, even if popular professional opinion has not" (Turri 2016, p.191).

[4]More exactly, Turri admits only that (1A*) *appears* questionable. For the sake of the argument, though, he concedes the dubious nature of (1A*). Despite believing that he can undermine reliabilism without relying on (1A*), he nevertheless adds that "... it seems to me that anything that can be reliably achieved can be unreliably achieved, so I would endorse the reading of 'achievements don't require reliable abilities' as a (necessarily true) universal generalization, even though such a reading isn't strictly necessary for my argument to pose a serious challenge to reliabilism" (Turri 2015a, p.534). As we shall see, my criticisms of Turri's main argument would also apply to this argument as well.

[5]While Turri explicitly states his argument (1A)–(3A), he never explicitly states (1A***)–(3A***). Nevertheless, he says this:

> ... if we understand 'achievements don't require reliable abilities' to express a tendency proposition, then line 2 [(2A***)] would still be very plausible. It would still be plausible because ... we would expect knowledge to fit the profile of achievement generally, unless we're given a special reason to think otherwise. And line 1 [(1A***)] would remain very plausible too, because of the long and varied list of types of achievement that don't require reliability. (Turri 2015a, p.534)

[6]Greco (2020) distinguishes roughly between "safe" and "unsafe" notions of achievement in a very brief discussion of Turri's examples.

[7]Thanks to an anonymous reviewer for pointing out a problem of my formulation of this issue.

[8]We could also look to a combination of slugging percentage and on base percentage (OPS). Many view this statistic as more important than batting average. As Houston Astros third baseman Alex Bregman puts it: "All the hitters care about now is OPS. We don't care about batting average. Some guys do, I guess. Batting average is an old stat that doesn't matter. It's OPS, runs created... Look at Mike Trout's numbers. There are guys that hit .340. Mike Trout is hitting, what, .300 on the dot? I'd rather have Mike Trout's numbers with all the walks and the damage than the guy who hits .340 [with a bunch of singles]. It's an OPS game" (Schoenfiled, 2019).

[9]Thanks to an anonymous reviewer for bringing this point to my attention.

[10]Perhaps here we should distinguish between ability and skill, as some have argued that Turri generally fails to make this distinction (see in particular McKinnon 2014, p.562).

[11]Thanks to an anonymous reviewer for comments that helped me make this point more precisely.

[12]Once again, an anonymous reviewer aided me in constructing more precise wording.

[13]The implausibility of the House show reminds me in some respects of the implausibility of the show *Murder She Wrote*. In that series, the mystery writer Jessica Fletcher is almost always around some strange murder. In one of the later shows, this implausibility is humorously noted by a character: "Jessica Fletcher is here and a murder's been committed. What are the odds?"

Acknowledgments

I wish to thank audiences at the Southeastern Epistemology Conference as well as the Alabama Philosophical Society for extensive discussions on much earlier versions of this paper. Ted Poston and Alex Whalen also gave me thoughtful comments on previous drafts. In addition, three anonymous reviewers for this journal provided extremely helpful feedback on these ideas. Finally, the University of South Alabama for reassigned time that helped me write this paper.