Augustana College

Augustana Digital Commons

Philosophy: Student Scholarship & Creative Works

Philosophy

Winter 2-23-2024

Anchoring the AAA Model

Luca James Barba
Augustana College, Rock Island Illinois, lucabarba20@augustana.edu

Follow this and additional works at: https://digitalcommons.augustana.edu/philstudent

Part of the Philosophy Commons

Augustana Digital Commons Citation

Barba, Luca James. "Anchoring the AAA Model" (2024). *Philosophy: Student Scholarship & Creative Works.*
https://digitalcommons.augustana.edu/philstudent/2

This Student Paper is brought to you for free and open access by the Philosophy at Augustana Digital Commons. It has been accepted for inclusion in Philosophy: Student Scholarship & Creative Works by an authorized administrator of Augustana Digital Commons. For more information, please contact digitalcommons@augustana.edu.

Luca James Barba ‘24 2021—2023

Anchoring the AAA Model

Part I: Animal Knowledge

Introduction

My goal in this paper is to expose two—but resolve one—major issues facing one of the foremost endeavors in the theory of knowledge: Virtue Epistemology. In Part 1 of this paper, I propose the epistemic criterion of "anchoring” as an addition to Ernest Sosa's AAA model of knowledge. It is a solution to an internal inconsistency that arises from Sosa's response to Duncan Pritchard's (2009) environmentalist luck critique of the AAA model that allows performances to spatiotemporally extend beyond their performers. By modifying the AAA model to the A4 (AAA + Anchoring) model, I can address Pritchard’s critique without disturbing traditional conceptualizations of the locations of mental performances. I end Part 1 by exploring three significant objections to anchoring.

In Part 2 of this paper, I explore issues that the AAA/A4 model has with accuracy and epistemic verification. While Sosa succeeds in granting someone who has insight into animal knowledge further access to reflective knowledge or the awareness of knowing, I argue that reflective knowledge cannot retroactively verify an apt belief’s truthfulness without circular reasoning. This issue arises from Sosa’s reliance on a Cartesian argument to support foundationalism in the face of the Pyrrhonian Problematic and Sextus’s Dark Room Scenario. The conclusion of Descartes’s conclusion that the treasure hunter in the Dark Room is required to gain awareness of the truth cannot be drawn from the premises Sosa has available.

Background

The Gettier Problem

The context for both Pritchard's and Sosa's work—and by extension, my own—is the monumental 1963 work by Edmund Gettier, "Is Justified True Belief Knowledge?". In this two-page paper, Gettier demonstrated that the multi-millennia-old "Knowledge is Justified True Belief" (JTB) model fails because a belief can be justified and true by pure coincidence. He explained this via an example of two men, Smith and Jones, seeking a job. The company's president assures Smith that Jones will get the job, and Smith counts ten coins in Jones's pocket. Given that any proposition deducible from a justified proposition is itself justified, Smith uses this information to form the justified belief that the person who gets the job has ten coins. Sure enough, this is true! Smith, not Jones, gets the job and finds ten coins in his pocket (1963, p. 59). Even though Smith's belief was justified and true, the connection between justification and truthfulness was arbitrary: the truthfulness came from the two men coincidentally having the same number of coins and not solely from Smith's learning and reasoning. That is unacceptable for a model of knowledge because the purpose of a model like this is to define a concept: just as an arbitrary definition fails to describe a word, an arbitrary model fails to illustrate a concept. This arbitrariness is now known as "Gettier luck," and it is why the JTB model definitively fails.

The AAA Model

In the wake of the "Gettier problem," which left epistemology without a plausible definition of knowledge, many epistemologists have offered replacements. One promising solution, the AAA model of knowledge, came from Ernest Sosa's 2007 lecture, "A Virtue Epistemology." Sosa uses a different perspective on the nature of beliefs to solve the Gettier problem. Instead of regarding

beliefs merely as attitudes or feelings, Sosa argues that beliefs, like many other human activities, are performances with aims (2007, loc. 224). This change in perspective is helpful because a successful belief would count as knowledge, meaning that one could solve the Gettier problem by answering the question, "What makes a performance successful?". For Sosa, this answer is outlined by the three A's of the AAA model:

> Accuracy: The performance must meet its goal; for a belief, this goal is being true.
>
> Adroitness: The performer must manifest some skill, ability, or competence; for a belief, this will typically be data-gathering or processing skills like seeing or hearing.[1]
>
> Aptness: The Adroitness of the performer must cause the belief to be accurate.

Because aptness entails accuracy and adroitness, he refers to successful performance as "apt," with apt beliefs constituting knowledge (2007, loci 212-215).

Sosa uses two examples of archers to explain the AAA model (2007, loci 212-217). In the first example, the performance is apt: seeing a target, accounting for wind speed, and using dexterity, an archer's arrow hits the bullseye. This performance is apt because it meets all three A's. The archer hits the desired target (accuracy). The archer uses various skills in the process (adroitness). And the performance is accurate because the archer is adroit (2007, loci 212-215). In the second example, the performance is accurate and adroit but fails to be apt. As in the first, the archer sees the target, accounts for wind speed, and uses his archery skills. The arrow even hits the bullseye. But unlike the first, the archer makes a slight miscalculation that would have caused the arrow to miss, and a freak gust of wind redirects the arrow to its target. Because of this last detail, the performance fails by

[1] For the sake of this paper, "skill," "ability," and "competence" are interchangeable. It is also worth noting that this criterion is why Sosa's field is called "Virtue Epistemology:" According to Aristotelian thought, virtues were qualities that allowed tools and people to achieve their goals—for example, sharpness is a virtue of a knife, and generosity is a virute of the soul. If the goal of a belief is to be true, there should be virtues that allow it to attain that goal.

AAA standards. It is accurate: the goal was to hit the bullseye, which happened. It is also adroit: the archer uses the same skills as before. But, it is not apt: the arrow hits the target because of a freak gust of wind, despite the archer's (mis)calculation (2007, loci 212-217). Simply put, apt performances are successful performances; not-apt performances are failed performances.

Importantly, this is also why the AAA model successfully responds to Gettier luck. Aptness saves the AAA model from Gettier luck because Gettier luck, according to Peter Unger, intervenes between the manifestation of skill and the moment of accuracy (1968, p. 159). In other words, Gettier luck occurs when the accuracy of a belief is not directly attributable to the believer's adroitness. Let us return to Smith and Jones to observe this. Recall that Smith held an accurate belief: the person who got the job did have ten coins in his pocket. Further, recall that Smith manifested adroitness: he counted the coins in Jones's pocket and competently communicated with the company president (1963, p. 59). But Smith's belief was not accurate directly because of his observation and reasoning; it was accurate because he and Jones had the same number of coins in their pockets, which had nothing to do with any manifestations of skill. Because of this subtlety, the AAA model holds that Smith did not have an apt belief—did not have knowledge—thus, the model is free from Gettier luck.

The Fake Barn Scenario

While this is a good start, this model has undesirable consequences in certain strange circumstances. Namely, the AAA model tolerates particular(ly) odd situations in which the likelihood of aptness is incredibly low. Consider Carl Ginet's "fake barn scenario" (Goldman 2000) for context. In the fake barn scenario, a woman finds herself in an unknown countryside. Also unknown to her is that this is fake barn country—while there are apparently many regular barns here, as there are in

any countryside, all but one of the barns are mere façades. The woman, looking around, happens to land her eyes on the one actual barn in the countryside. Observing this barn, she draws the true belief that there is a barn in the countryside (Goldman 2000).

Although the woman develops a true belief in this scenario, many epistemologists consider such a belief inadequate for knowledge because it falls under the category of environmental luck: while her senses may have been completely functional, the prevalence of fake barns—the likelihood that the environment would have deceived those senses—is so high that the true belief ought to be attributed to a certain luck rather than skill. One such epistemologist with this perspective is Duncan Pritchard, who argues that the AAA model is susceptible to *environmental* luck. Pritchard explains that the woman in Ginet's fake barn scenario meets AAA standards. Her belief is accurate, as it is true that there is a real barn. Her belief is adroit, as the woman used her senses. And her belief is apt, as those senses generated the accurate belief. So, the woman in Fake Barn Country clearly knows that there is a barn per the AAA model (2009, p. 411).

Seeing this and other consequences of the AAA model that Pritchard demonstrated, Sosa argues that environmental luck is epistemologically acceptable but offers modifications to the AAA model. First, knowledge ceases to be binary: something as simple as an automatic door can have *some* brute knowledge of those who step in its sensor. This knowledge is inferior to the knowledge that the woman has of the barn, which itself is lesser than the knowledge that I have of my bulldog's name (Sosa 2009, p. 431). As I understand it, Sosa argues that all beliefs exist along a spectrum of epistemic success, and any clear-cut distinction between knowledge and true belief—assuming that the belief is apt—is linguistic rather than metaphysical. Second, Sosa explains that abilities can interact serially and simultaneously. For example, one can know that there are seagulls when seeing the silhouettes of birds in the distance and hearing cawing sounds, even if neither experience alone

would have allowed one to determine which species one is observing. As for serial abilities, one can piggyback off the epistemic labor of predecessors, such as one's teachers, to learn things that one could never discover by oneself. Third and finally, "aptness" does not merely mean that a skill caused a true belief to form; it means that a skill caused the belief to be true (2009, p. 433).

These changes to the AAA model dissolve the issue of environmental luck, at least for our friend in the countryside. Through the first modification, the model can account for the diversity of robustness seen across epistemic performances. Because circumstances like environmental luck can detract from the qualitative—rather than the binary—success of a true belief, the woman's belief of there being a barn counts as AAA knowledge while not being considered equal to obviously superior epistemic achievements, such as my knowledge of my bulldog's name. So, her belief is partially discredited, even if it still counts as knowledge. Through the second modification, the AAA model can explain why her knowledge is inferior to mine. The skills used in the fake barn scenario could have easily misled the woman, so her epistemic performance is flimsier than those conducted in optimal conditions. But my regular exposure to my family discussing little Vinnie over the past few years is unlikely to have misled me. Finally, through the third modification—which is more of a clarification—Sosa explains why the model does not allow for Gettier luck. In Gettier scenarios like the one with Smith and Jones, the use of skill does cause the formation of a belief; however, unnoticed environmental circumstances intervene outside of any manifestation of competence and cause that belief to be accurate. So, Gettier beliefs fail to be apt.

Testimonial Knowledge

While these modifications resolve critical issues for the AAA model, they also cause new ones. These new issues pertain to the boundaries of beliefs, which are exemplified by testimonial

knowledge. Promoted by thinkers like Jennifer Lackey, testimonial knowledge refers to knowledge acquired through communication between agents (2007). It is an everyday phenomenon. Seeing diagrams in an instruction manual, reading a biology textbook, and hearing an emergency alert broadcast are all instances of gaining information through testimony. In the AAA model, Sosa accounts for testimonial knowledge through credibility-based trust and listening skills. Specifically, one can rely on the credit that others have earned through consistently offering truthful information to aptly form testimonial beliefs (2009, p. 432). Instruction manuals, biology textbooks, and emergency alerts are all examples of trustworthy sources of information; con artists, antiquated medical manuscripts, and yellow journalists are not. In his response to Pritchard, Sosa acknowledges that the bulk of the epistemic burden could fall on someone other than the believer, such as in a student-teacher interaction, in which the student's epistemic burden is limited to moments of listening while the teacher's entails years of study in the relevant field and its pedagogy (Sosa 2009, p. 435).

Making Knowledge Make Sense

AAA Issues With Testimonial Beliefs

When abilities can serially and simultaneously interact in the context of testimony, however, the boundaries of one's actions become unacceptably strange. Presupposing that (1) one's beliefs cannot exist beyond oneself, and (2) a belief's adroitness is a component of that belief, then one's beliefs can now spatially extend beyond oneself. Recalling that beliefs are activities that contain thoughts and abilities in the AAA model, the location of a belief is limited to where one embodies the relevant abilities and has the relevant thoughts. But this is no longer the case when abilities serially and simultaneously interact in testimony, as one's belief can extend to where someone else embodied an

ability. Consider an everyday, Lackeyian scenario where a middle school biology teacher teaches her students about the theory of evolution: as the students acquire knowledge through her, they leverage her abilities in their epistemic performances (Lackey 2007). As their budding beliefs would serially include her abilities, it seems that those beliefs would spatially extend to the teacher. And insofar as beliefs are actions, it would logically follow that the students' actions are partially done by the teacher. Further, if the teacher were a Creationist who denied evolution but still taught it, it would follow that the belief in evolution would extend to someone who did not believe in evolution.

This line of reasoning extends to the second issue that arises with testimony: temporal boundaries. In other words, Sosa's modifications of the AAA model in response to Pritchard make it so that the children's acts of believing in evolution precede their lifespans. As Darwin and Wallace were the first individuals to discover evolution (ignoring Empedocles for brevity), it would follow that the children's beliefs about evolution could extend beyond the temporal boundaries of their lives, all the way back to those two. After all, the children testimonially inherited the abilities of Darwin and Wallace through their teacher, as those abilities are part of the causal chain of abilities that formulated the modern conceptualization of evolution. Even if one denied that Darwin, Wallace, and the teacher were actors in the children's beliefs, it would still follow that their beliefs spatiotemporally extended to include these individuals, as their abilities are considered parts of the children's beliefs per the modified AAA model.

Interestingly, Sosa enters a dilemma, as he cannot deny this spatiotemporal extension of beliefs without re-opening the door to environmental luck. If abilities do not serially link beyond oneself, then environmental luck can prevail through testimony. Returning to the fake barn scenario, assume that the woman had built up an ample supply of epistemic credit over the years: she is a well-known expert in barn architecture and surely a reliable source for barn-related information. Had a curious

child asked her about the barn she saw, she could surely teach him about the barn in barn country. But that child would have qualitatively exceptional knowledge of the barn, as no abilities contained within his epistemic performance were conducted in sub-optimal conditions. The woman was the one who exercised those weak abilities, and she did not pass those on to the child.

This issue extends to any form of luck and even allows for weird lies and so-called "bullshit" (in Frankfurt's 1986 sense of the word) to create knowledge. Because the testifier's epistemic success becomes irrelevant when the ability chain ends with the listener, an unsuccessful testifier can instill false positives for knowledge in listeners. Had the woman seen a fake barn, mistook it for a real one, and passed that information onto the child, the child would have met the conditions for knowledge. He would have embodied skills—including listening to a creditable (although mistaken) barnologist—that caused him to produce an accurate belief. So, the child would certainly know of the barn per AAA standards, but he would have done so through Gettier luck! Clearly, this is unacceptable for epistemology.

Anchoring: A Solution

While there is more than one way to address this, I suggest adopting the A4 model, which adds anchoring to the modified AAA model. Instead of only linking abilities, as Sosa argues, anchoring entails linking whole performances. Each atomic performance still needs to meet the original three A's; however, when a particularly long series of abilities occurs with several intuitively discrete aims, the series counts as several performances rather than one. Conveniently, these performances will interact similarly to how Sosa links abilities together. Depending on the context, sub-performances will compose or precede super-performances, and latter performances succeed only if prior performances did. Furthermore, "super-performance" has two related definitions: in the first, it is a

standalone performance that relies on some other performance to succeed; in the second, it is an over-arching set of performances that come together to facilitate an end.

There are two ways that sub-performances can relate to each other and super-performances: chaining and weaving. It can be hard (and oftentimes irrelevant) to distinguish the two, but this is the difference: Performance A is chained to Performance B if and only if Performance B influences the aptness of Performance A; Performances A and B are woven with respect to Performance C if and only if neither influences the other's aptness, but the success of either one could contribute to the accuracy of Performance C.

To exemplify this, I borrow J L Mackie's scenario involving one of my fellow arsonists burning a house (2000). In this scenario, burning down the house is a super-performance because it is the all-encompassing performance of a set of interrelated but spatiotemporally discrete (sub-)performances. The arsonist may be diligent: stripping the insulation off many wires in the attic, leaving a lit candle in the basement, and turning on the stovetop with kitchen grease nearby. These activities are woven with respect to the super-performance of arson: each could contribute to the arsonist's success in igniting a house fire, although one success will not contribute to the success of another. On the other hand, the super-performance is directly chained to each of these sub-performances: at least one must succeed for the super-performance to succeed. This also clarifies how super-performances meet the anchoring condition in practice. A successful super-performance will follow from a continuous chain of successful sub-performances. Regardless of how these sub-performances interrelate, there must be some chain of performances, unbroken by failure, for the super-performance to qualify as a success.

Three integral caveats are visible once one considers the subtleties of chaining and weaving. First, there can be broken chains within a super-performance—there can be failed

sub-performances—while the super-performance succeeds. Second, sub-performances can sometimes be included in super-performances that deviate from the goals of the sub-performances. For example, an incompetent general contributes to his opponent's success by organizing an incompetent army. And third, there can be redundant chains—more than one completely successful lineage of sub-performances—but the performance only counts as one success. Suppose the arsonist got his way, and the house caught fire. But it only caught fire from the kitchen stove. The chains from the super-performance to the stripped wires and lit candle would be broken by the failure of those two actions to result in ignition. In this way, the wire stripping would be unsuccessful, as would the candle lighting—but the super-performance of arson would still succeed because of lighting the kitchen stove. Conversely, had all of those sub-performances been successful, the super-performance would still have succeeded, even though it only necessitated the success of one. The old AAA model could not clearly account for this nuance.

With the anchoring condition, the A4 model allows for the success of a belief to remain tethered to aptness while not opening the door to epistemic performances preceding their believers or knowledge arising from Gettier luck through testimony. How the A4 model achieves the first success is quite simple: when one develops a belief through testimony, the testifier's testifying counts as a distinct performance from the learner's development of a true belief. In this way, beliefs have boundaries that do not extend beyond believers. The categorization of performances into super- and sub-performances then solves the issue of false positives for knowledge through testimony. When instilling knowledge in the listener becomes a shared super-performance between the listener and the testifier, in which the listener's epistemic performances include listening and processing, and the testifier's epistemic performances include gathering information, processing it, and then passing it on through testimony, it is easy to point (from an omniscient perspective) at any failures in that chain and know that super-performances—even if resulting in true beliefs—fail.

Objections

Can Super-Performances Be Susceptible To Luck?

A worthwhile concern is whether this merely puts virtue epistemology back at large-scale square one. Can super-performances as a whole fall to Gettier or environmental luck? Consider a scenario in which all of humanity is put into Matrix-esque holding tanks, where every human epistemic performance takes place in a simulation (Wachowskis 1999). But at the same time, the alien inhabitants have created mindless androids who mirror the humans' appearances and actions perfectly in real-time, such that the material world and the Matrix are perpetually observably identical. In such a scenario, every human performance and super-performance would occur in the Matrix but yield true beliefs about the material world.[2]

The consequences of anchoring here vary upon the presuppositions and conclusions that the entrapped humans draw. *Prima facie*, anchoring still allows the entrapped humans to learn some things from within the matrix. "I perceive beige walls," "I feel pain," and "I want to go to bed" are all conclusions that one could successfully draw, as one concludes nothing about the external world. (Notice how all of these are "I experience" statements.) But success becomes strange in cases like "It is nice outside." If "outside" means "the perceivable world," then the belief succeeds because the matrix is the perceivable world. Technically, "outside" could mean "the material world" and potentially succeed, as there is a non-arbitrary, perfect alignment between the material world and the matrix. But such a claim would likely rely on the inference that "I am observing the material world;" in which case, that belief would fail.

[2] This challenge was suggested to me by my advisor, Deke Gould.

What's The Deal With Memory?

The second objection focuses on the limitations of anchoring. When there is a great distance between a successful sub-performance and a later super-performance, the *continuing* aptness—if such a thing exists—of a sub-performance can be attributable to luck rather than skill; anchoring does not account for this and, therefore, still allows for luck. To exemplify this, I would like to return to my bulldog, Vinnie. Vinnie has an odd habit of burying his bones in my backyard just to retrieve them months later. But consider a scenario in which Vinnie was cryogenically frozen for a millennium, much like Fry in Matt Groening's *Futurama* (1999).

Further, assume that the Qu, an alien race, had conquered the solar system at this time, aggressively reshaping much of Earth's geography and genetically modifying mankind into elaborate playthings (Kösemen 2008). Vinnie—now awake and characteristically indifferent to the Qu—merrily navigates to his buried bone, digs it up, and starts gnawing away. In this scenario, Vinnie's super-performance meets A4 standards. It is accurate: he ended up with his bone. It is adroit: Vinnie used active recall. It is apt: Vinnie retrieved his bone because of his recollection. And it is anchored: Vinnie's recollection was the product of his earlier, successful performance of memorizing where he had buried his bone. So, Vinnie should be considered a very good boy by A4 standards.

But despite his success per A4 model standards, Vinnie was obviously lucky that his bone remained where he had left it and that he could navigate to its location with a long-outdated understanding of the world. This raises many questions about the place of time and memory in virtue epistemology. First, at what point is it lucky that one's memories reflect current reality—further, are memories ever to be trusted in this way? Second, what even is a memory—is it a sustained belief over time or a reservoir of old information? Could memories be their own

multi-step performances that entail memorizing, retaining, and recalling, or are they abilities? Might there be different types of memories?

Questions like these are worth answering. Anchoring would allow for memories to be easily described as super-performances. For example, remembering where one's house is includes observing one's house, correctly inferring that one's house is real, observing and correctly noting the surrounding area, weaving it all together for navigation, and then inferring that the data remains accurate. But, I am not qualified to make definitive statements about memory's workings, so I cannot answer this here.

Part II: Reflective Knowledge

Introduction

In Part II of my paper, I examine how anchoring impacts Sosa's methodology for attaining reflective knowledge of aptness's epistemological value. I argue that Sosa makes a subtle misstep in the argument for justifying how he can produce the AAA model in the face of several skeptical threats. Specifically, he conflates precision for accuracy, which allows him to make an apparently cohesive defense against the Pyrrhonnian Problematic, the Criterial Problem, and Sextus's Dark Room Scenario using Descartes's *scientia*. I argue that this conflation offers Sosa a dubious entrance into reflective knowledge of the aptness's reliability as it results in circular reasoning but not circular justification. Finally, I respond to an objection.

Background

The Three Challenges To Epistemology

After the lecture in which he argues for the AAA model, Sosa provides another lecture on how the model can withstand skeptical threats, specifically those that challenge the ability to create epistemological models. The three main threats that Sosa responds to are Roderick Chisholm's Criterial Problem, the Pyrrhonian Problematic, and Sextus's Dark Room scenario. They all revolve around a Criterion: a standard that a belief could meet to constitute knowledge (Conee 2009, p. 419). Considering this in the context of Peter Unger's "knowledge is a true belief had in the right kind of way," a Criterion is a right kind of way (1968). Each challenge, in one way or another, holds that a Criterion is out of grasp.

The first and most recent challenge is Roderick Chisholm's Criterial Problem (1973). The Criterial Problem sets up a paradox whereby knowledge of a Criterion could not come without already knowing a Criterion. It takes the following form: to epistemically progress, one can either analyze what one already knows to discover a Criterion, or one can have a Criterion to discover what one knows. The paradox is that one cannot conduct any form of analysis to discover a Criterion without requiring access to fully justified knowledge, which itself requires the very Criterion the undertaking is meant to find. Because of this Criterial Problem, Chisholm holds, one has no access to knowledge (1973).

The next issue Sosa faces is the Pyrrhonian Problematic. Developed by Pyrrho of Elis (ca. 370-272), the Problematic holds that a belief's rationale always fails when it depends on an arbitrary or unjustified criterion. Effectively, this requires justified beliefs to be justified by other justified beliefs (Sosa 2009, loc. 1080). Sosa writes that the Pyrrhonian Problematic results in three possible

structures for epistemic practice. First, there is a denial of the Problematic in which one argues that some beliefs or Criteria are justified without relying upon anything else. This is the foundationalist approach. Second, the infinitist approach agrees that each justified belief stands on the backs of other justified beliefs and holds that there are infinite regresses in epistemology. Third, there is the option where one agrees with Pyrrho in holding that each justified belief stands on the backs of other justified beliefs. However, one also holds that the regresses can circle back in on themselves, becoming self-supporting. These are the coherentists (Sosa 2007, loc. 1082).

Finally, Sosa has to respond to Sextus's Dark Room scenario. That scenario questions whether, even with limited information available, one can discover a Criterion. It takes on the following form:

> Let us imagine taht some people are looking for gold in a dark room full of treasures… [N]one of them will be persuaded that he has hit upon the gold even if he has, in fact, hit upon it. In the same way, the crowd of philosophers has come into the world, as into a vast house, in search of truth. But it is reasonable that the man who grasps the truth should doubt whether he has been successful (Sextus Empiricus 1914, p. 52).

The predicament here is that even if one finds the gold—even if one has true beliefs in the right kind of way—the fact that one currently has no way of becoming aware of that means that one is barred from finding that out. While there may be knowledge or a Criterion, the current inaccessibility of a Criterion prevents thinkers from ever discovering one (Sosa 2007, loc. 1082).

A Brief Introduction To Reflectivity

Sosa provides interrelated answers to each of these questions. He does this in three ways: He makes the AAA model bi-layer, relies on foundationalism, and takes a gradual approach to knowledge acquisition. The bi-layer structure of the AAA model is that it contains animal knowledge and reflective knowledge. As a reminder, this paper has only discussed Sosa's lowest tier of

knowledge so far—animal knowledge—which Sosa defines as "apt belief" (Sosa 2007, loc. 228). That sort of knowledge is quasi-externalist insofar as one can have apt beliefs without any awareness of their aptness (Sos 2009, loc. 343). Sosa defines the higher tier of knowledge as "apt belief aptly noted," and it serves a confirmational role in the AAA model (Sosa 2009, loc. 305).

Reflective knowledge occurs when one has attained apt belief with a level of clarity and careful consideration such that one can be certain that one's belief is successful (Sosa 2009, loc. 315). The theoretical value of reflectivity is it allows one to "break into" awareness of a true belief in everyday situations. So, when I see a keyboard under my fingers, I have animal knowledge that there is a keyboard. But when I also note that the lighting is good, that I can easily see the keyboard, and that I am not descending into madness, I have achieved the "aptly noted" part of "apt belief aptly noted." I can now say that I reflectively know there is a keyboard under my fingers.

How Reflectivity Avoids These Challenges

Reflectivity avoids Chisholm's Criterial Problem in a subtle but impressive way. According to Sosa, the Criterial Problem only restricts Criteria from discovering themselves; it does not restrict one form of knowledge from discovering another. For example, if one already knows certain truths per Criterion A, one need not cite Criterion B to use those truths to discover Criterion B. In this way, one can use animal knowledge to illuminate reflective knowledge, bypassing the Criterial Problem (Sosa 2007, p. 131). Rather than being two purely distinct Criteria, however, reflective knowledge and animal knowledge share a Criterion: aptness. But this is not an issue for reflective knowledge, as reflective knowledge has an additional Criterion: aptly noting. So, Sosa argues, relying on truths yielded from animal knowledge can grant access to a new Criterion, which is aptly noting

(Sosa 2007, p. 131). Other thinkers, such as Earl Conee, have come out in support of Sosa's solution to Chisholm's Criteria Problem (Conee 2009, p. 418).

While this may address the Criterial Problem, it raises the question of where one's support for animal knowledge is to come from. Even if animal knowledge can provide insight into reflective knowledge, thinkers intuitively need support from something else to justify their reliance on animal knowledge. This is the critique seen in the Pyrrhonian Problematic: to have a fully justified belief in anything, one's belief must be justified by beliefs that are themselves fully justified (Sosa 2007, loc. 1080). Of the three available options for dealing with this—foundationalism, infinitism, and coherentism—Sosa opts for foundationalism, arguing that it is the only way of avoiding the Pyrrhonian Problematic (Sosa 2007, loc. 1108).

This is no free lunch for the AAA model, as Sosa goes on by responding to a core issue facing foundationalism in this context: why does an observational belief not require justification? Although Sosa consults several thinkers, including Lawrence BonJour and Wilfrid Sellars, this critique of foundationalism culminates in Richard Rorty's claim that causation ought not to be confused with justification. From Rorty, there are two important points: (1) simply because one has beliefs, and those beliefs are clearly caused by something, it does not follow that those beliefs are *justified* by that cause; and (2) "nothing counts as justification unless by reference to what we already accept, and there is no way to get outside our beliefs and our language so as to find some test other than coherence" (Sosa 2007, loc. 1100; LePore 1986, pp. 307-320).

Interestingly, Sosa outright states that he agrees with Rorty (Sosa 2007, loc. 1102). He states, however, that this is no issue for the AAA model, and turns to Descartes to explain how.

> The fact that an atheist can be "clearly aware that the three angles of a triangle are equal to two right angles" is something that I do not dispute. But I maintain that this awareness of his

> [*cognitionem*] is not true knowledge [*scientia*], since no act of awareness that can be rendered doubtful seems fit to be called knowledge [*scientia*]. Now since we are supposing that this individual is an atheist, he cannot be certain that he is not being deceived on matters which seem to him to be very evident (as I fully explained). And although this doubt may not occur to him, it can still crop up if someone else raises the point or if he looks into the matter himself. So he will never be free of this doubt until he acknowledges that God exists (Descartes 1984, p. 101).

In this example, Descartes formulates knowledge to have two tiers, a lower and a higher. Even with the threat of deception, Descartes holds, the atheist can possess *cognitio*[3]—the lower tier of knowledge that lacks any awareness of success. He argues that the atheist cannot attain *scientia* because, in doubting God, the atheist is unable to prove that any of his reasoning is free from deception—knowing that one is free from deception is entry-level knowledge.

As far as he is concerned, this nuance is all Sosa needs to dodge Rorty's threat to foundationalism and the Pyrrhonian Problematic. Drawing from the *cognitio*/*scientia* distinction, Sosa argues that animal knowledge can succeed without any awareness of its possessor, much like the atheist (Sosa 2007, loc. 1147). Definitionally, this is because there is no "aptly noted" requirement in animal knowledge; the apt believer need not be certain of his belief's aptness. But more technically, this is because the primary difference between the atheist and Descartes is not which tools they use but whether they can confirm that those tools are reliable. So, as long as the atheist uses the right tools, he can still achieve limited knowledge. In this way, the AAA model takes on a semi-externalist stance, evading Rorty's concerns that assume that *all* forms of justification require awareness of success for success.

Interestingly, this is also how Sosa responds to Sextus's Dark Room scenario. Rather than Descartes's long-rejected approach that the only way to *scientia* is belief in God, Sosa argues that one can work with animal knowledge to discover reflective knowledge. He agrees with Sextus that the

[3] *Nota bene*: I am not qualified to conjugate *cognitio* or *scientia* for grammatical case and will not do it in this paper.

practically blind men rummaging for gold are unable to attain any awareness of success with a mere "thermometer's reaction" to their environment (Sosa 2007, loc. 1162). However, Sosa argues that the men could, with enough data gathered through observation, gain reflective knowledge of which objects are gold and which are not (Sosa 2007, loc. 1160). Insofar as simple aptness requires no internal awareness of that aptness, all that is left for them to adopt the presupposition that their sense data is worthwhile and build up:

> If we persist in such reasoning, nevertheless, enough pieces may eventually come together into a view of ourselves and our place in the universe that is sufficiently comprehensive and coherent to raise us above the level of mere *cognitio* and into the realm of higher, reflective, enlightened knowledge, or *scientia*. (Sosa 2007, loc. 116).

In this way, the treasure hunters can break from the externalist animal knowledge into the internalist reflective knowledge while avoiding all major skeptical threats (Sosa 2207, loc. 1160).

A Subtle Misstep

An argument as subtle and cohesive as this is undoubtedly praiseworthy; nonetheless, there is an issue. Simply put, Sosa uses animal knowledge to discover reflective knowledge in the face of the Criterial Problem; however, he uses reflective knowledge to give one the epistemological insight (but not justification) to use observation to yield animal knowledge. This is a mistake in the context of anchoring, as someone attempting to move from a comprehensive and cohesive worldview to the reflective knowledge that that worldview is true, would be using the presupposition "My senses sometimes yield truth" to conclude "My senses sometimes yield truth." This issue does not exist for a Cartesian model of higher knowledge, which Sosa relies on to make his argument, but does exist for Sosa's own model of reflective knowledge.

As a reminder, Sosa and Conee are correct in holding that having a bi-layer model of knowledge allows one to escape *Chisholm*'s Criterial Problem—in the case of the AAA model, one uses aptness (Criterion A) to discover aptly noted aptness (Criterion B) (Sosa 2007, p. 131; Conee 2009, p. 418). But then arises the Pyrrhonian Problematic: how can one *trust* aptness without some other justification?

Sosa's misstep begins when he applies Descartes's argument to the Dark Room Scenario. Recall that Descartes identifies lower-tier knowledge with insight from higher-tier knowledge. When he argues that the atheist mathematician possesses *cognitio* of geometry, Descartes writes from the perspective of having already attained *scientia* himself, then reflects on what separates his knowledge from that of the atheist. It is not because of Descartes's *scientia* that the atheist can achieve *cognitio*, but it is because of someone's *scientia* that anyone can be aware that the atheist's beliefs count as *cognitio*. Because of that, the atheist is necessarily oblivious to the extent of his *cognitio*.

Here, one sees why the treasure hunters searching for gold in Sextus's Dark Room are not analogous to Descartes's atheist searching for truth. In Sextus's Dark Room, the treasure hunters entered the room already having *scientia* that (1) there is gold and (2) gold has observable properties. In other words, there is no presupposition needed that adroitness can yield truth. In contrast, Descartes's atheist lacks awareness that there is animal knowledge or how to get it. He cannot take for granted any epistemological right-of-way; he simply has data-yielding practices. Full-heartedly adopting the presupposition that aptness (or, in practice, adroitness) is the path to animal knowledge, even if externally justified, would be a leap of faith as far as the atheist can tell. For this reason, framing the Dark Room scenario from the perspective of Descartes's atheist would look fundamentally different. Essentially, Sextus's treasure hunter and the atheist treasure hunter both have various data-yielding practices available, but only Sextus's hunter has *scientia* that any are reliably

truth-yielding: He walked into that room knowing that gold exists in there and that it is detectable through the senses. The atheist treasure hunter, on the other hand, may have a presupposition that his senses are truth-yielding and a hope that it is correct.

Implicit in this distinction is a serious mistake that Sosa makes in his formulation of the AAA model, and this mistake will cost the atheist the ability to become his own Descartes. Namely, Sosa conflates precision and accuracy. In archery, precision is when an arrow hits a bullseye; accuracy is when an arrow lands near recently-launched arrows. Precision is absolute, while accuracy is relative. In the context of beliefs produced from abilities, precision would be a measure of how similar a belief is to the truth, while accuracy would be a measure of how similar a belief is to another belief.

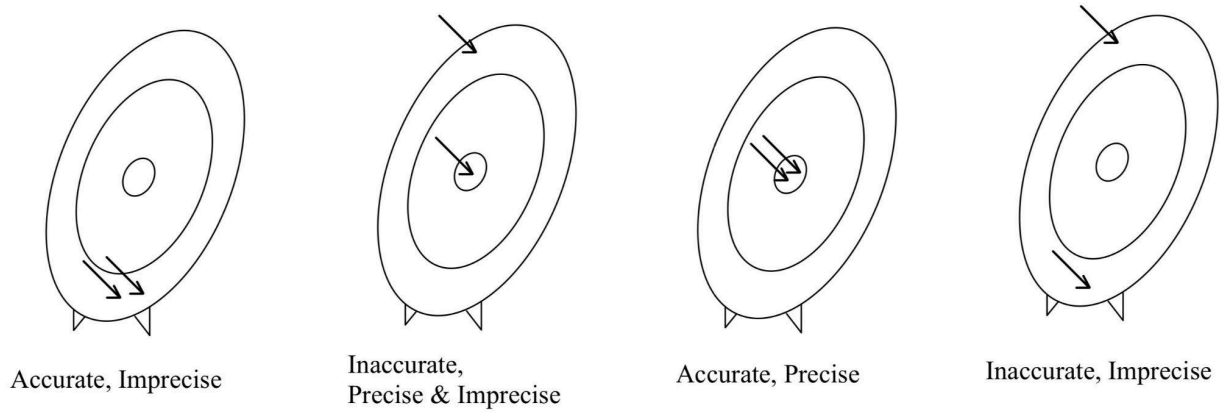

Considering his definition of accuracy, "the belief is true" or "the arrow hits its target," in the context of this diagram, one notices that Sosa is not actually discussing accuracy; he is discussing precision (2007, loc. 212).

This is a non-trivial issue because there is a clear way to argue that human senses can yield accurate beliefs. Present 100 healthy but blindfolded anglophones a coffee, and they will all perceive the scent of coffee: their arrows all land in the same place, and their beliefs are mutually accurate.

When one seeks to validate a scientific finding, peers scrutinize the study's methodology and aim to replicate its data. Even on a large scale, people frequently (but certainly not always) adopt similar beliefs when uncovering new phenomena: consider how English has adapted over the years to account for new species, technologies, and social trends so that anglophones can easily index and discuss them. When Sosa writes that having a "comprehensive and cohesive" worldview allows one to gain access to reflective knowledge, then, it strongly implies that this reflectivity is a conditioned awareness of which beliefs are accurate rather than which are precise (2007, loc. 1160).

This leads to two main issues. First, note that linking accuracy and precision entails an explanatory circularity for Sosa's reflectivity. In analyzing the accuracy/precision gap in the context of anchoring, one notices a limitation of observation: observation is merely the experience of certain qualities—redness, roundness, etc.; concluding that those qualities reflect the external world is an inference. Presupposing that inference to succeed is the exact presupposition Sosa requires the treasure hunters—in other words, us—to make to progress from unaware, animal knowledge to aware, reflective knowledge (2007, loc. 1160). But the thing is, no matter how cohesive or comprehensive it is, the performance of creating one's worldview will inevitably be anchored in the presupposition that one's senses can sometimes yield truth. If one were to look at this process as an argument with premises and conclusions linked through anchored practices, one would notice that the men in the dark room would be assuming the premise "My senses are sometimes truth-yielding" to derive the conclusion "My senses are sometimes truth-yielding." While potentially justified from an externalist perspective, it is circular reasoning as far as the men can reflectively tell. This is the ultimate challenge posed by the accuracy/precision gap: to prove that very accurate beliefs are precise, one must assume that very accurate beliefs are precise.

Second, this is also where one sees that Sosa's reflective knowledge, thought to be analogous to Descartes's *scientia*, cannot do the heavy lifting that Sosa needs: proving that aptness is truth-yielding. This issue arises because Sosa's reflective knowledge rises from unverified premises, while Descartes's *scientia* does not. If a treasure hunter were to find consistent practices to detect and conceptualize his experiences through trial and error, he would be able to produce a comprehensive and cohesive worldview. He could even learn how to reflectivley determine when a manifestation of skill is unlikely to be reliable, as he will notice that similar, past performances produced inconsistent or low-resolution outcomes. However, he still presupposes accuracy and adroitness to be standards for success. He will need to go back and verify aptness's status as a Criterion before he can be aware that accuracy aligns with precision. This is fundamentally distinct from Descartes's conceptualization of *scientia*, which by design has no reason to retroactively validate any of its founding principles. As far as Descartes was concerned (although he was wrong), the formulation of *scientia* entailed only undeniable presuppositions (Descartes 2018, p. 10). Had Descartes succeeded, there would have been no need to retroactively very anything; there would be no circularity issue seen in Sosa's reflectivity.

Objections

There is one glaring objection to my argument that draws back to the original presentation of the AAA model. That objection is Sosa's requirement for basis-relative safety, a direct countermeasure to skeptical threats like my accuracy/precision gap. A belief is safe when one holds it only if it is likely to be true (Sosa 2007, loc. 235). In the context of Sosa's virtue epistemology, which treats beliefs as performances, a belief is safe when it is unlikely to have failed (Sosa 2007, loc. 233). There are plenty of examples of everyday beliefs that fail to be safe—Sosa provides one of a hypochondriac who experiences a minor discomfort but confuses it for excruciating pain (Sosa

2007, loc. 245). That is why he instead supports basis-relative safety: a belief that is basis-relative safe has a basis that one would likely have only if true (Sosa 2007, loc. 251).

Basis-relative safety is a direct countermeasure to my accuracy/precision criticism because it evades the Cartesian Demon threat (Sosa 2007, loc. 256). With the intervention of a Cartesian Demon, all of one's sensory experiences are called into question. Until one can disprove the Demon's existence, according to the skeptic, one cannot trust any sort of knowledge. Basis-relative safety avoids this in two ways. First, beliefs that there is Demonic intervention in perception lack any sort of clear and present basis and thus fail as knowledge (Sosa 2007, 225). Second, beliefs that there are no Demonic interventions are very based and thus can succeed as knowledge (Sosa 2007, loc. 264).

My response to this objection is that, while it may succeed in preventing the AAA model from counting Demonic perceptions as knowledge, it still fails in the same way that relying on Descartes did before. Beginning from a perspective of complete ignorance like the atheist treasure hunter, one lacks insight that would indicate which bases would be present only if true or which performances would be unlikely to fail. In other words, the treasure hunter has no idea what constitutes a substantive base, so basis-relative sensitivity is of no practical use from his perspective.

Conclusion

My goals in this paper have been to provide an important solution, anchoring, and explore its consequences for Ernest Sosa's AAA model. In Part I, I introduced anchoring and demonstrated that it could resolve an abnormality that arises when abilities within a single belief are linked across time and individuals through testimony. That abnormality was that a belief could spatiotemporally extend beyond the individual believer, including to doubters who simply contributed to it. I ended

that Part by exploring two objections, one concerned with anchored performances being susceptible to Gettier luck and another concerned with how memories relate to anchoring. In Part II, I presented Sosa's explanation of how one is to move from epistemic blindness to reflective knowledge. I demonstrated how his reliance on Descartes and an accuracy/precision conflation results in an issue of circular reasoning when a hypothetical man attempting to discover aptness from reflective knowledge seeks to explain his discovery.

Citations and References

Chisholm, Roderick. "The Problem of the Criterion." Marquette U P, 1973.

Descartes, René. *The Philosophical Writings of Descartes*. Translated by John Cottingham, Robert Stoothoff, and Dugald Murdoch, vol. 2, Cambridge U P, 1984.

Frankfurt, Harry G. "On Bullshit." Raritan Quarterly Review, vol. 6, no. 2, 1986, pp. 81-100.

Gettier, Edmund. "Is Justified True Belief Knowledge." *Epistemology: An Anthology*. Edited by Ernest Sosa and Jaegwon Kim, 1st ed., Blackwell Publishing, 2000, pp. 58-59. Print.

Goldman, Alvin. "Epistemic Folkways and Scientific Epistemology." *Epistemology: An Anthology*. Edited by Ernest Sosa and Jaegwon Kim, 1st ed., Blackwell Publishing, 2000, pp. 438-444. Print.

Kösemen, C. M. *All Tomorrows*. Flesk Publications, 2008.

Lackey, Jennifer. "Why We Don't Deserve Credit for Everything We Know." *Synthese*, 2007, pp. 158, 345-362.

Lepore, Ernest. "A Coherence Theory of Truth and Knowledge." *Truth and interpretation: Perspectives on the philosophy of Donald Davidson*. Blackwell, 1986, pp. 307-320.

Mackie, JL. "Causes and Conditions." *Epistemology: An Anthology*. Edited by Ernest Sosa and Jaegwon Kim, 1st ed., Blackwell Publishing, 2000. Print

The Matrix. Directed by Lilly Wachowski and Lana Wachowski, performances by Keanu Reeves, Laurence Fishburne, Carrie-Anne Moss, Hugo Weaving, and Joe Pantoliano, Warner Bros., 1999.

Pritchard, Duncan. "Apt Performance and Epistemic Value." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, vol. 143, no. 3, 2009, pp. 407–416. *JSTOR*, www.jstor.org/stable/27734414. Accessed 21 Mar. 2021. Digital.

Sextus Empiricus. *Against the Mathematicians*. Edited by H. Mutschniann, 1914.

Sosa, Ernest. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford U P, Inc, 2007. Digital.

———. "Replies to Brown, Pritchard, and Conee." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 2009.

"Space Pilot 3000." *Futurama*, created by Matt Groening, season 1, episode 1, 20th Century Fox Television, Mar. 28, 1999.

Unger, Peter. "An Analysis of Factual Knowledge." *Journal of Philosophy*, 1968.