Knowledge in Degrees

<u>Abstract</u>

In contemporary epistemology, it is predominantly held that propositional knowledge is a binary, all-or-nothing notion, rather than a graded notion. In this paper, I challenge the predominant view. I begin with some motivations for an account of graded knowledge. I then consider various notions pertaining to belief and justification, and I explore their potential roles in an account of graded knowledge. Finally, I use the various notions of belief and justification to develop a formula for determining a subject's degree of knowledge for a given proposition.

1. Motivations

A key motivation for developing an account of graded knowledge is summarized by Stephen Hetherington, in his book *How to Know: A Practicalist Conception of Knowledge*:

> Voluminous philosophical discussion exists, concerning what *kind* of justification is required [for knowledge]…. But how *strong* must knowledge's justificatory component be? On this, there is a startling silence…. [A]ll that is usually suggested is that a true belief's being knowledge involves its being *well* justified (or other words to similar effect). And what, precisely, does that mean? We are yet to be told…. The situation is not one where we find a plethora of individual epistemologists vigorously defending their respective delineations of knowledge's justificatory boundary – and then disagreeing with each other over its location, thus failing to reach a consensus. Instead, almost all epistemologists are simply silent on these details (2006, p. 41-42).

What might explain this "startlingly silence" regarding the problem? One possibility is that there just isn't any definitive answer to the question, "How *strong* must knowledge's justificatory component be?". BonJour argues for something along these lines when he writes:

> [J]ust *how* probable must it be in light of the reasons or justification available that the proposition is true in order for it to be adequately justified to count as knowledge? … [I]t is simply unclear what sort of basis or rationale there might be for fixing this level of justification in a nonarbitrary way (2002/2009, p. 39).

Thus, if there is no clear answer to the threshold problem, it is perhaps unsurprising that epistemologists tend to remain silent on the issue. However, this leaves things in a fairly uncomfortable position. For, if there can be no definitive response to the threshold problem, and the threshold problem threatens the standard conception of fallibilist knowledge, then what lasting value does the fallibilist conception of knowledge have?

Accordingly, BonJour goes on to suggest that the threshold problem calls into question "whether the concept of knowledge has the importance that is often attributed to it: how important could it be (and why) that the strength of one's reason or justification for a claim is above rather than below a line that cannot be clearly and nonarbitrarily defined?" (BonJour 2002/2009; also see Hintikka 2006). Quine takes a similar stance, writing:

> The notion of knowledge is beset … by a less subtle difficulty: a vagueness of boundary…. We do better to accept the word 'know' on par with 'big', as a matter of degree. It applies only to true beliefs, and only to pretty firm ones, but just how firm or certain they have to be is a question, like how big something has to be to qualify as big…. [T]he best we can do is give up the notion of knowledge as a bad job and make do rather with its separate ingredients. We can still speak of belief as true, and of one belief as firmer or more certain, to the believer's mind than another (1987).

Is this the only way forward? In particular, might there be some systematic way in which the "separate ingredients" fit together and thereby reveal how the notion of knowledge is in fact useful? I submit that there is. In particular, if knowledge is graded as some have suggested (e.g., Weiler 1965; Hetherington 2001; Sosa 2007), and if the normatively interesting notions involved in knowledge include not only the particular ingredients (such as, degrees of belief and degrees of justification) but also the normative relations between these ingredients (e.g., the relation between one's degree of belief and degree of justification), then this could motivate an account on which a subject's degree of knowledge for some proposition is determined by the subject's degree of justification, the subject's degree of belief, and relationship between the subject's degree of justification and degree of belief.

To further motivate an account of graded knowledge, consider the following scenario. Suppose that Charles is a meteorologist living in Boston. After a careful, systematic study of the weather conditions in Boston, Charles' evidence clearly suggests that there is an 85% chance of snow in Boston this evening. Thus, Charles' degree of belief that it will snow in Boston this evening equals 0.85.

Micah, who knows little about meteorology, lives in Chicago. As he is leaving his apartment to go to work, his roommate says to him, "You should wear winter boots. A man in the elevator just said there's an 85% chance of snow this evening." Micah has no further evidence that there is an 85% chance of snow in Chicago this evening, but he also does not have any evidence to the contrary. However, unlike Charles, whose degree of belief (0.85) accords with his evidence that there is an 85% chance of snow in Boston this evening, Micah rather rashly forms a degree of belief of 0.99 that it will snow in Chicago this evening.

Supposing that it will in fact snow in both Boston and Chicago this evening, what should we conclude about Charles' and Micah's knowledge regarding the respective propositions? There certainly seems to be some difference between their overall epistemic states: for one, the overall epistemic state of Charles seems to be better than that of Micah. For, although the evidence that is available to Charles indicates the same chance of snow in Boston (85%) as the evidence available to Micah regarding Chicago (85%), Charles' evidence for the proposition <there is an 85% chance of snow in Boston this evening> is clearly stronger than Micah's evidence for the proposition <there is an 85% chance of snow in Chicago this evening>. Moreover, while Charles' degree of belief accords with his evidence, Micah's degree of belief goes beyond what is indicated by his evidence.

Standard binary accounts of knowledge seem unequipped to tell us much about how such variables might contribute to one's overall epistemic state, and thus seem unable to tell us much about any overall epistemic difference between the states of Charles and Micah. By contrast, I submit that an account of graded knowledge has potential in addressing such issues.

2. A virtue-theoretic approach to graded knowledge

What should an account of graded knowledge look like, and in particular, how should it relate to the notions of belief and justification? One possible answer comes from work in virtue epistemology, i.e., a general approach to epistemology in which the notion of intellectual virtue plays a central role. Contemporary interest in virtue epistemology can be largely traced back to a series of papers written by Ernest Sosa in the 1980s and early 1990s (Sosa 1980, 1985, 1988, 1991). According to Sosa's account, which finds its most mature and systematic expression in his subsequent work (Sosa 2007, 2009, 2011), knowledge is best understood as a type of skillful performance.

In short, Sosa (2007) argues that performances may be evaluated in terms of three features: "accuracy: reaching the aim; adroitness: manifesting skill or competence; and aptness: reaching the aim *through* the adroitness manifest". Sosa's predominant example is that of an archer's shot. Three questions are relevant to evaluating the shot. First, is the shot accurate: does it hit the target? Second, is the shot adroit: does it manifest the archer's skill or competence? Third, is the shot apt: is its accuracy a result of the archer's competence? These three questions correspond to what Sosa calls the "AAA structure" of performances (accuracy, adroitness, aptness), and he argues that this structure is well suited for the evaluation of beliefs. He writes, "We can distinguish between a belief's accuracy, i.e., its truth; its adroitness, i.e., its manifesting epistemic virtue or competence; and its aptness, i.e., its being true *because* competent" (Sosa 2007).

Given Sosa's characterization of knowledge as a type of skillful performance, the question arises as to whether knowledge, like other types of skillful performances, can come in degrees. Sosa (2007) suggests that it can, writing, "[K]nowledge is a matter of degree, in various respects," and he notes that this accords with the idea of "knowledge as apt belief, since aptness of belief admits degrees". Nevertheless, Sosa develops his account primarily in binary terms. Regarding accuracy, he writes, "we can…assess how accurate a shot it is, how close to the bull's-eye, [but] we here put degrees aside, in favor of the on/off question: whether it hits the target or not" (Sosa 2007). Similarly, regarding adroitness, he notes, "Skill [adroitness] too comes in degrees, but here again we focus on the on/off question: whether it manifests relevant skill or not, whether it is or is not adroit" (Sosa 2007). And lastly, regarding aptness, he states, "Aptness is a matter of degree even beyond the degrees imported by its constitutive adroitness and accuracy, for a performance is apt only if its success is *sufficiently* attributable to the performer's competence" (Sosa 2007).

While Sosa puts degrees aside in favor of the on/off question, I put the on/off question aside in favor of degrees. Namely, I focus on how the general AAA structure of knowledge proposed by Sosa provides a useful groundwork from which to build an account of graded knowledge.

3. Optimal degree of belief

In accordance with Sosa's AAA model of knowledge, the following proposal may serve as a provisional starting point for an account of graded knowledge:

> *Prov1* The degree to which a subject *S* knows some true proposition *P* is an increasing function of: (i) the degree to which *S* is *justified* (*competent*) in believing *P*, and (ii) the degree to which *S believes P*.

While *Prov1* has some intuitive pull to it, there are a few caveats that need to be made if it is to inform an account of graded knowledge. First, I take it for granted that if *P* is false, then *S*'s degree of knowing *P* equals 0. Second, in accordance with the gradability of a belief's *accuracy*, I aim to construct an account of graded knowledge that is tied to the notion of graded belief (rather than binary belief).[1] Given this, the current wording of (i) within *Prov1*—i.e., "the degree to which *S* is *justified* (*competent*) in believing *P*"—lacks a clear reference. In particular, I cannot expect any answer to the question: "To what degree is *S* justified (competent) in believing *P*?" Rather, I need to specify a particular degree of belief, for which I can then ask about the corresponding degree of justification.

---

[1] I use phrases such as "*S*'s degree of believing *P*" to refer to the strength of *S*'s tendency to think as though *P*.

What particular degree of belief should I specify in order that I may then determine the corresponding degree of justification? I do not want to use *S*'s actual degree of belief, since that would result in an account on which the justification (competence) condition entails the belief condition. To illustrate, consider a parallel with the justified true belief (JTB) account of knowledge. The justification condition in the JTB account does not require that one first consider whether *S* actually believes *P* before it can be determined whether *S* is justified in believing *P*. If this were the case, then the justification condition would entail the belief condition, and thus the belief condition would be superfluous, which it is not.

So, if I do not want to use *S*'s actual degree of belief here, what degree of belief should be used in order to determine the corresponding degree of justification? One way to address this issue is to introduce the notion of an *optimal degree of belief* (abbreviated B_O). In short, a subject's B_O is the subject's most highly justified degree of belief, relative to all of the other degrees of belief that the subject might have.

To illustrate how a subject's B_O might differ for various propositions, suppose that Omar is a college student and is currently waiting in a lecture hall for his biology class to start. During the wait, two things occur: first, Omar reads in his biology textbook that cells vary widely in the number of their mitochondria, and second, he happens to overhear one of his classmates, whom he knows little about, say to someone else, "I think the lecture this coming Friday will be canceled." Supposing that both of the relevant propositions are true—i.e., that cells vary widely in the number of their mitochondria, and that the lecture this coming Friday will indeed be canceled—one might suppose that Omar's B_O for the former proposition is around 0.9 (that is, out of all the degrees of belief that Omar might have, he is most justified in believing *P* to a degree of 0.9), while Omar's B_O for the latter proposition is a fair amount lower, perhaps around 0.6; since Omar's reasons suggest a higher probability for the former proposition than the latter.

However, it is important to distinguish between a subject's B_O (i.e., the subject's degree of belief with the highest degree of justification, relative to all other degrees of belief that the subject might have) and a subject's *perfectly (100%) justified degree of belief*. In particular, to say that a subject's B_O equals *n* for some proposition *P*, does not imply that the subject is *perfectly* (100%) justified in believing (to a degree of *n*) that *P*; or equivalently, using a scale of 0 to 1, it is *not* to say that the subject is justified (to a degree of 1) in believing (to a degree of *n*) that *P*. To see why, consider two parallel scenarios.

Sasha takes a coin out of his pocket and has no reason to think it's unfair. Thus, his B_O for the proposition <the coin will land heads> is 0.5 (that is, out of all the degrees of belief that Sasha might have for the proposition <the coin will land heads>, there is no better degree of belief that he could have than one of 0.5).

Oscar likewise takes a coin out of his pocket, but after doing so, he performs careful experiments measuring the coin's physical properties to ensure that the coin is fair (which it is). Thus, Oscar's B_O regarding the proposition <the coin will land heads> also equals 0.5. However, given that Oscar has much stronger evidence (relative to Sasha) that the probability of the coin landing heads equals 0.5, Oscar's degree of justification in believing (to a degree of 0.5) that *P* holds is greater than Sasha's degree of justification in believing (to a degree of 0.5) that *P* holds. So, while Oscar and Sasha have the same B_O for *P*, the degree to which Oscar's B_O is justified is greater than that of Sasha's. Or to put it another way, for both Oscar and Sasha, the degree of belief with the highest degree of justification (relative to all of the other possible degrees of belief that Oscar and Sasha might have) equals 0.5; but for Oscar, his highest degree of justification (i.e., the degree to which he is justified in his B_O) is significantly higher than Sasha's highest degree of justification. For instance, perhaps the degree to which Oscar is justified in his B_O is, say, 0.95, whereas the degree to which Sasha is justified in his B_O is a bit lower, say, 0.7.

4. The degree to which S is justified in a given B_O

This raises the question of how a subject's degree of justification (competence) in having a B_O for some proposition should be used to formulate the justification (competence) condition (from Sosa's AAA account) in order that it might be better adapted to an account of graded knowledge. Recall that the notion of a subject's optimal degree of belief (i.e., B_O) was introduced so that we would have a particular degree of belief (namely, a subject's B_O for some proposition *P*) for which we could then ask about the corresponding degree of justification (i.e., the degree to which the subject is justified in having a particular B_O for *P*) that corresponds to a degree of knowledge. It may at first be thought that it is only the *degree of justification* (which corresponds to the subject's B_O) that is needed in the reformulation of the justification (competence) condition, where the notion of B_O merely plays an auxiliary role in helping one determine the corresponding degree of justification. To see why this idea is problematic, consider two cases:

> Sophie is a meteorologist and is currently living in El Paso. After a careful, systematic study of the weather conditions in El Paso, her evidence provides strong support *against* the proposition <it will snow in El Paso this evening>. Accordingly, Sophie's B_O for this proposition is

> very low (say, 0.07). Moreover, given Sophie's training as a meteorologist, and her careful study of the weather conditions, the degree to which she is justified in having this B_O is quite high (say, 0.9).

Now, contrast this scenario with the following scenario:

> Nicolas, who knows very little about meteorology, currently lives in Boston. As he is leaving his apartment to go to work, his roommate says to him, "You should wear some winter boots. I just heard a man in the elevator say that there's an 85% chance of snow this evening." Assuming that Nicolas has no evidence to the contrary, his B_O for the proposition <it will snow in Boston this evening> equals 0.85. That is, given the reasons available to Nicolas (however limited they might be), there is no degree of belief that is better supported by Nicolas's reasons than a degree of belief equal to 0.85. But what is the degree to which Nicolas is justified in believing (to a degree of 0.85) that it will snow in Boston this evening (that is, the degree to which he is justified in his B_O)? The answer here seems to be that he has a moderate degree of justification (say, 0.65).

So, in sum, Sophie's B_O for the proposition <it will snow in El Paso this evening> is 0.07, and the degree to which she is justified in having this B_O is 0.9. By contrast, Nicolas's B_O for the proposition <it will snow in Boston this evening> is 0.85, and the degree to which he is justified in this B_O is 0.65.

Consider why these scenarios are relevant to an account of graded knowledge. First, suppose that both of the propositions under consideration are true: that is, against all odds, and to the astonishment of meteorologists, it will in fact snow in El Paso this evening; and, in accordance with the suggestion made by Nicolas's roommate, it will indeed snow in Boston this evening. If the only relevant notion to the justification (competence) condition, as revised for an account of graded knowledge, is the degree to which a subject's B_O for some proposition is justified, then this leads to a problem. To illustrate, let JB_O stand for the degree to which S's B_O for some proposition is justified. Now, assume (for the purpose of reductio) that S's degree of knowledge is an increasing function of S's JB_O (without any concern for the value of S's B_O, other than its use in determining *S*'s JB_O). Then, Sophie might know the proposition <it will snow in El Paso this evening> to a significantly higher degree than Nicolas knows the proposition <it will snow in Boston this evening>. But this seems untenable. For Sophie has strong evidence that it will *not* snow in El Paso this evening, while Nicolas has moderately good evidence that it *will* snow in Boston this evening. Thus, how could Sophie's degree of knowledge for the relevant proposition be greater than Nicolas's degree of knowledge for the relevant proposition?

A similar problem arises if one takes the view that *S*'s degree of knowledge is an increasing function of *S*'s B_O (without any concern for the value of *S*'s JB_O). For instance, suppose that Charles is a meteorologist who lives in Boston. After a careful, systematic study of the weather conditions, Charles' evidence clearly indicates that there is a 95% chance of snow in Boston this evening. Thus, Charles' B_O for the proposition <it will snow in Boston this evening> equals 0.95. And given his scientific training and thorough analysis of the weather conditions, the degree to which Charles is justified in having this B_O is very high, say, 0.9.

Micah, by contrast, knows very little about meteorology. While eating dinner in a restaurant, he happens to overhear someone say that there's a 95% chance of snow in Boston this evening. Given that Micah has no evidence to the contrary, we may suppose that his B_O is 0.95. However, since Micah has a significant lack of information about both the speaker and the context in which the claim is uttered, the degree to which he is justified in having this B_O is quite low, say, 0.3.

Thus, supposing that it will in fact snow in Boston this evening, it seems that (ceteris paribus) Charles has a significantly greater degree of knowledge regarding the relevant proposition than does Micah. This leads to a problem for the view that the only relevant notion to the revised justification (competence) condition is a subject's B_O. For, according to such a view, Charles and Micah have an *equal* degree of knowledge that it will snow in Boston this evening.

5. Degree of justificatory weight

In light of the above arguments, I propose a revised version of the justification (competence) condition that incorporates both B_O and JB_O. Specifically, I propose that a subject's (*S*'s) degree of knowledge is an increasing function of *S*'s *degree of justificatory weight* (abbreviated J_W) with respect to *P*. The value of *S*'s J_W depends on both the value of *S*'s B_O and the value of *S*'s JB_O. Before giving a formula to determine a subject's J_W regarding some proposition, I should offer a few remarks about how a subject's B_O and JB_O relate to the subject's degree of justificatory weight (i.e., J_W), and thus how they relate to the subject's degree of knowledge.

First, I propose that when *S*'s B_O is greater than 0.5 (i.e., when *S*'s reasons suggest that *P* is probable), then a higher JB_O results in a higher degree of knowledge (and thus, for the purpose of developing the notion of J_W, it results in a higher value for J_W). This idea is illustrated in the above scenario in which both Charles and Micah have a B_O that equals 0.95 for the proposition <it will snow in Boston this evening>; and while Charles' JB_O equals 0.9, Micah's JB_O only equals 0.3. It was then suggested that (ceteris paribus) Charles has a higher degree of knowledge than Micah regarding the proposition <it will snow in Boston this evening>. This accords with the idea that when B_O is greater than 0.5, a higher JB_O results in a higher degree of knowledge.

Next, consider a situation in which *S*'s B_O is *less* than 0.5. Here, something quite different occurs. Specifically, when *S*'s B_O is less than 0.5, a higher JB_O seems to result in a *lower* degree of knowledge (and thus, for the purpose of developing the notion of J_W, it should result in a *lower* value for J_W). To elaborate, if *S*'s B_O is less than 0.5 (and hence *S*'s reasons suggest that *P* is *improbable*), then the higher the degree is to which *S* is justified in having this B_O, the less *S*'s degree of knowledge is that *P* (and thus the less *S*'s J_W should be for *P*).

For example, consider a case in which Jesse, a meteorologist, has very strong reasons *against* the proposition <it will snow in Los Angeles this evening>. Specifically, Jesse's B_O for this proposition is 0.15, and her JB_O (i.e., the degree to which she is justified in having this B_O) is quite high (say, 0.9). Further suppose that Caitlin also has a B_O of 0.15 for the proposition <it will snow in Los Angeles this evening>. But unlike Jesse, Caitlin's B_O is not due to careful, systematic studies, but instead results merely from overhearing a passerby mention to someone else that there's a 15% chance of snow this evening (while having little information about the speaker or the context of the conversation). Accordingly, Caitlin's JB_O is quite low, say, 0.3.

Given that Jesse has much *stronger* reasons (relative to Caitlin's) *against* the proposition in question, this seems to suggest that Jesse is at a disadvantage (relative to Caitlin) in obtaining a high degree of knowledge that it will snow in Los Angeles this evening. Suppose, for instance, that against all odds, it actually happens to snow in Los Angeles this evening. Then, if there is any difference between the degree to which Jesse knows this proposition and the degree to which Caitlin knows it, it seems that (ceteris paribus) Jamie's degree of knowledge is *less* than that of Caitlin's, since Jamie has *stronger* reasons *against* the proposition.[2]

Finally, there is one last condition to consider before providing a formula for J_W. Namely, it seems that the greater the distance that a subject's B_O is from 0.5, the greater the effect that the subject's JB_O has on the subject's degree of knowledge (and accordingly, on the subject's J_W). To illustrate, suppose that Dave has a tendency to walk around parks and give coins to people, truthfully informing them that the coin is unfair and that there's a 95% chance that it will land heads. When Dave gives Imran a coin, Imran receives the coin and merely accepts Dave's claim that there is a 95% chance that it will land heads. Thus, assuming that Imran has no particular reason to doubt the claim, his B_O for the proposition <the coin will land heads> is 0.95.

Sharva likewise receives a coin from Dave, but after doing so, Sharva performs careful experiments measuring the coin's physical properties to ensure that the coin is in fact weighted such that there is a 95% chance that it will land heads (which there is). Thus, Sharva's B_O regarding the proposition <the coin will land heads> equals 0.95 (the same as that of Imran). However, given that Sharva has much stronger evidence (relative to Imran) that the probability of the coin landing heads equals 0.95, Sharva's degree of justification in having this B_O (of 0.95) is much greater than that of Imran's. In particular, given the rigorous tests that Sharva performs on the coins physical properties, one might say that the degree to which Sharva is justified in having this B_O is 0.9 (that is, Sharva's JB_O is 0.9). By contrast, Imran's B_O (which, like Sharva's, equals 0.95) is

[2] I realize, however, that it is not wholly implausible to think that Jesse and Caitlin might just have the *same* degree of knowledge (i.e., a degree of knowledge that equals 0). For, one might hold that if a subject's B_O is not greater than 0.5, then the subject's degree of knowledge equals 0, regardless of the value of other variables. The manner in which this alternative view would affect the formula for J_W is considered in Footnote 3.

merely due to him hearing Dave's claim. Since Imran knows little about Dave, other than his odd behavior of handing out unfair coins, the degree to which Imran is justified in his B_O is quite low (say, 0.3).

Supposing that when Imran and Sharva flip their respective coins they will both land heads, it seems that (ceteris paribus) Sharva has a significantly greater degree of knowledge regarding the proposition <the coin will land heads> than does Imran.

Now, compare this case with one in which Sharva and Imran both have a B_O that equals 0.52 (rather than 0.95). Suppose, for instance, that the scenario is such that Dave gives out coins while truthfully informing people that there is a 52% chance that it will land heads. Moreover, as in the previous case, Sharva performs careful studies of the coin's physical properties to ensure that it is weighed in such a way that there is a 52% chance that it will land heads (which there is). So, in this case, Sharva's B_O equals 0.52, while the degree to which he is justified in this B_O still equals 0.9.

Imran likewise has a B_O that equals 0.52. And as before, his B_O is merely due to hearing Dave claim that there is a 52% chance that the coin will land heads. Thus, the degree to which Imran is justified in this B_O is again quite low, say, 0.3.

To summarize, Sharva has very strong evidence that there's a 52% chance that the coin will land heads, while Imran has fairly weak evidence that there's a 52% chance that the coin will land heads. Supposing that they both flip their respective coins and that both coins will in fact land heads, which of the two (Sharva or Imran) has the greater degree of knowledge that the coin will land heads? While Sharva may have a greater degree of knowledge regarding P than that had by Imran, it seems that the difference in their respective degrees of knowledge is less pronounced in this case than in the case in which both Sharva and Imran have a B_O that equals 0.95. Thus, in accordance with the above suggestion, the greater the distance of a subject's B_O from 0.5, the greater the effect of the subject's JB_O on the subject's degree of knowledge (and hence, on the subject's J_W).

5.1. Determining a subject's J_W

In light of the above discussion, the formula for J_W should satisfy the following three conditions:

(c1) When B_O is greater than 0.5, a higher JB_O results in a higher J_W.

(c2) When B_O is less than 0.5, a higher JB_O results in a lower J_W.

(c3) The greater the distance of B_O from 0.5, the greater the effect JB_O has on the value of J_W.

To construct a formula for *S*'s justificatory weight (J_W) that satisfies (c1) – (c3), begin by considering a case in which *S*'s optimal degree of belief (i.e., B_O) equals 0.8:

(i) $B_O = 0.8$.

Then, suppose that the degree to which *S* is *justified* in this B_O (i.e., *S'*s JB_O) equals 0.2:

(ii) $JB_O = 0.2$.

Given that *S* is 20% justified in her B_O of 0.8, it seems that at least 20% of *S*'s B_O should count towards her justificatory weight (that is, at least 20% of 0.8 should count towards J_W):

(iii) J_W has a value of at least 0.16 ($0.2 \cdot 0.8 = 0.16$).

Now consider the remaining value of B_O (i.e., the value determined by, $B_O - B_O \cdot JB_O$). This value equals 0.64 ($0.8 - 0.16 = 0.64$). What percentage of this remaining value (if any) should count towards J_W? Initially, it may seem that none of B_O's remaining value should count towards J_W. But this leads to a problem. Suppose that none of B_O's remaining value counts towards J_W. Then, as indicated in (iii), J_W equals 0.16. Hence, the value of J_W is the product of B_O and JB_O ($J_W = B_O \cdot JB_O$). While this may initially seem plausible for cases in which $B_O > 0.5$, it does not work for cases in which $B_O < 0.5$. As noted in condition (c2), when B_O is less than 0.5, a higher JB_O should result in a lower J_W; but given the formula, $J_W = B_O \cdot JB_O$, a higher JB_O would result in a higher (rather than a lower) J_W. So perhaps, instead, we can use the following two-part formula to define J_W:

$$Jw = \begin{cases} B_O \cdot JB_O, & \text{if } B_O > 0.5 \\ B_O \cdot (1 - JB_O), & \text{otherwise} \end{cases}$$

While this looks more promising, it's still not satisfactory. Suppose that *S* has very strong evidence that there's a 50.01% chance that it will rain, while *S** has very strong evidence that there's a 49.99% chance that it will rain. Thus, *S*'s B_O equals 0.5001, while *S**'s B_O equals 0.4999. Further suppose that both *S* and *S** have a JB_O of 0.9. Then, given the above formula, *S*'s J_W equals 0.45009, while *S**'s J_W only equals 0.04999. This is problematic, since it seems untenable to think that, with JB_O held fixed (at 0.9 for both subjects), such a small difference between *S*'s B_O and *S**'s B_O (a difference of 0.0002) would result in such a drastic difference between *S*'s J_W and *S**'s J_W (a difference of 0.4001). The point becomes even more vivid when comparing this case with a case in which the difference between *S*'s B_O and *S**'s B_O still equals 0.0002, but in which the space between *S*'s B_O and *S**'s B_O falls just *above* 0.5 rather than spanning *across* 0.5. For instance, suppose that *S*'s B_O is 0.5001 while *S**'s B_O is 0.5003. Then, assuming that *S* and *S** both still have a JB_O of 0.9, the difference between *S*'s B_O and *S**'s B_O still equals 0.0002 (as before). But now there is no longer such a drastic difference between S's J_W and S^*'s J_W (i.e., the difference equals 0.00018 rather than 0.4001).

One way to resolve this issue is suggested by condition (c3), above. It states: the greater the distance of B_O from 0.5, the greater the impact of JB_O on the value of J_W. Accordingly, when $B_O > 0.5$, the percentage of B_O's remaining value that should count towards J_W may be calculated as follows:

$$(B_O - (B_O \cdot JB_O)) \cdot 2(1 - B_O)$$

And thus J_W may be calculated by adding this value to $B_O \cdot JB_O$, as follows:

$$Jw = B_O \cdot JB_O + \big(B_O - (B_O \cdot JB_O)\big) \cdot 2(1 - B_O), \qquad \text{if } B_O > 0.5$$

On the other hand, when $B_O < 0.5$, a higher JB_O should result in a *lower* (rather than a higher) J_W. So, the corresponding formula for cases in which B_O is less than (or equal to) 0.5 is:

$$Jw = B_O \cdot (1 - JB_O) + \big(B_O - \big(B_O \cdot (1 - JB_O)\big)\big) \cdot 2B_O, \qquad \text{if } B_O \leq 0.5$$

These two formulas may be simplified into the following two-part formula:[3]

$$Jw = \begin{cases} B_O \cdot (1 - (1 - JB_O) \cdot 2|0.5 - B_O|), & \text{if } B_O > 0.5 \\ B_O \cdot (1 - JB_O \cdot 2|0.5 - B_O|), & \text{otherwise} \end{cases}$$

To see how this two-part formula works, reconsider the case in which S's B_O equals 0.5001, while S^*'s B_O equals 0.4999, and in which both subjects have a JB_O of 0.9. The

[3] If one prefers the view suggested in Footnote 2 (i.e., that if a subject's B_O is not greater than 0.5, then the subject's degree of knowledge equals 0), then J_W could instead be formulated as:

$$Jw = \begin{cases} B_O \cdot (1 - (1 - JB_O) \cdot 2|0.5 - B_O|), & \text{if } B_O > 0.5 \\ 0, & \text{otherwise} \end{cases}$$

problem with the earlier proposed formula was that such a small difference in the value of B_O (a difference of 0.0002) resulted in a very large difference in the value of J_W (a difference of 0.4001). But with this new formula, the problem is avoided. Regarding *S*, for instance, the formula, $B_O \cdot (1 - (1 - JB_O) \cdot 2|0.5 - B_O|)$, equals 0.50009. Regarding *S**, on the other hand, the formula, $B_O \cdot 1 - (1 - JB_O) \cdot 2|0.5 - B_O|$, equals 0.49989. So, rather than having such a drastic difference in the value of J_W (i.e., 0.4001), the difference in the value of B_O is now appropriately much lower (i.e., 0.50008 – 0.49989 = 0.00019).

Next, compare the difference between two more cases: one in which a subject's B_O equals 0.52, while JB_O equals 0.3; and another in which a subject's B_O equals 0.52, while JB_O equals 0.9. Regarding the first situation (in which B_O equals 0.52, and JB_O equals 0.3), the formula, $B_O \cdot (1 - (1 - JB_O) \cdot 2|0.5 - B_O|)$, equals 0.5054. Hence, J_W equals 0.5054 (0.0146 less than B_O). Now, consider the second situation (in which JB_O equals 0.9, and B_O still equals 0.52). Here, the formula, $B_O \cdot (1 - (1 - JB_O) \cdot 2|0.5 - B_O|)$, equals 0.5179. So, in this situation, J_W equals 0.5179 (0.0021 less than B_O). Thus, in brief, when B_O is held fixed at 0.52, the difference between (a) the degree to which a JB_O with a value of 0.3 affects the value of J_W, and (b) the degree to which a JB_O with a value of 0.9 affects the value of J_W, amounts to a difference of 0.0125 (0.5179 – 0.5054).

But consider how this same method of calculation applies when B_O equals 0.95 (rather than 0.52). By following the same steps, it can be seen that when B_O equals 0.95, the difference between (a) the degree to which a JB_O with a value of 0.3 affects the value of J_W, and (b) the degree to which a JB_O with a value of 0.9 affects the value of J_W, amounts to a much larger difference: i.e., one of 0.513,[4] rather than 0.0125 (as is the case when B_O equals 0.52). Thus, in accordance with what the above proposal, i.e., (c3), the greater the distance a subject's B_O from 0.5, the greater the effect of the subject's JB_O on the subject's J_W.[5]

The values calculated here match the values in the two parallel scenarios with Sharva and Imran (Section 5). Recall that, in the first scenario, Sharva and Imran both have a B_O that equals 0.95 (and Sharva's JB_O equals 0.9, while Imran's equals 0.3), In the second scenario, Sharva and Imran both have a B_O that equals 0.52 (and Sharva's JB_O equals 0.9, while Imran's equals 0.3). Thus, given the above calculations, the difference between Sharva's J_W and Imran's J_W amounts to 0.513 in the first scenario but only 0.0125 in the second scenario. Assuming that *S*'s degree of knowledge is an increasing function of *S*'s J_W, these calculations accord with the suggestion that the difference between Sharva's degree of knowledge and Imran's degree of knowledge is greater in the first scenario than in the second.

I now turn to consider cases in which B_O is less than 0.5. How might the formula apply to the case of Jesse and Caitlin (Section 5)? Recall that Jesse is a meteorologist and has very strong reasons suggesting that there is a 15% chance of snow in Los Angeles this evening; thus, she has strong reasons *against* the proposition <it will snow in Los Angeles this evening>. By contrast, Caitlin, who is not a meteorologist, has fairly weak reasons suggesting that there is a 15% chance of snow in Los Angeles this evening. Accordingly, one might say that Jesse and Caitlin both have a low B_O, i.e., 0.15. But

---

[4] When $B_O = 0.95$, it follows that: If $JB_O = 0.3$, then $J_W = 0.3515$; and if $JB_O = 0.9$, then $J_W = 0.8645$; thus, the difference is calculated at 0.513 ($0.8645 - 0.3515 = 0.513$).

[5] Also note that if $B_O = 0.5$, then $J_W = 0.5$. In other words, when $B_O = 0.5$, the value of JB_O has no influence on the value of J_W. This results from the proposal that the greater the distance between a subject's B_O and 0.5, the greater the effect of the subject's JB_O on the subject's J_W.

whereas Jesse's JB_O is quite high (say, 0.9), Caitlin's JB_O is quite low (say, 0.3). Now, supposing that against all odds it actually happens to snow in Los Angeles this evening, it seems that (ceteris paribus) Jamie's degree of knowledge is lower than that of Caitlin's; since, relative to Caitlin, Jamie has *stronger* evidence *against* the proposition.[6] This fits well with the above formula for cases in which B_O is less than or equal to 0.5: namely, the formula, $B_O \cdot (1 - JB_O \cdot 2|0.5 - B_O|)$. To illustrate, when Jesse's B_O and JB_O are entered into the formula, the resulting value is 0.0555.[7] By contrast, when Caitlin's B_O and JB_O are entered into the formula, the resulting value is 0.1185.[8] Thus, granting that one's degrees of knowledge is an increasing function of J_W, these values are in agreement with the above suggestion that (ceteris paribus) Jamie's degree of knowledge is lower than that of Caitlin's.

With the notion of a subject's J_W now in hand, I reformulate *Prov1* as follows:

> *Prov2* The degree to which a subject *S* knows some true proposition *P* is an increasing function of: (i) the value of *S*'s justificatory weight (i.e., J_W), and (ii) the degree to which *S believes P*.

6. Degree of doxastic weight

The next task is to consider the second condition of the provisional account. Specifically, consider the wording of (ii) within *Prov2*—i.e., "the degree to which *S believes P*". There seems to be some plausibility to the claim that if *P* is true, then *S*'s degree of knowing *P* is an increasing function of *S*'s degree of actual belief. However, this is only true for certain cases regarding *S*'s degree of actual belief: namely, those in which *S*'s degree of actual belief (abbreviated B_A) is less than *S*'s B_O. To illustrate,

[6] For an alternative viewpoint, and a corresponding alternative formulation of J_W, see Footnotes 2 and 3.

[7] $0.15 \cdot (1 - 0.9 \cdot 2|0.5 - 0.15|) = 0.15 \cdot 0.37 = 0.555$

[8] $0.15 \cdot (1 - 0.3 \cdot 2|0.5 - 0.15|) = 0.15 \cdot 0.79 = 0.1185$

suppose that the value of Omar's B_O for the proposition <the lecture this coming Friday will be canceled> is 0.6, as posited in Section 3. Then, if Omar's B_A is already at (or above) 0.6, it seems untenable to think that Omar could further increase his degree of knowledge by increasing his B_A to some unduly high level, say, 0.9. In fact, if Omar were to do this, it would seem to *decrease* rather than increase his degree of knowledge. By contrast, it *does* seem tenable to think that if Omar's B_A is less than his B_O (say, his B_A is 0.3, while his B_O is still 0.6), then he *could* increase his degree of knowledge by increasing his B_A to 0.6.

To incorporate this point, I introduce the notion of a subject's *degree of doxastic success* (abbreviated *Ds*) for some proposition. In short, a subject's *Ds* is an increasing function of both: the extent to which B_A reaches B_O, and the extent to which B_A does not surpass B_O. This idea may be expressed more precisely, as follows:[9]

$$Ds = \max(B_O - |B_O - B_A|, 0)$$

First, consider the subformulas. The subformula, $|B_O - B_A|$, is used to calculate the absolute difference (i.e., the distance) between *S*'s degree of actual belief (B_A) and *S*'s optimal degree of belief (B_O). This value is then subtracted from B_O using the subformula, $B_O - |B_O - B_A|$. In essence, the resulting value represents the proximity of B_A to B_O, on a scale of 0 to B_O. So, if B_A is less than or equal to B_O, the resulting value is just B_A.[10] If B_A is greater than B_O (indicating that there's an excess of B_A relative to B_O), then the resulting value is determined by subtracting this excess amount from B_O. Therefore, if Omar's B_O for the proposition <the lecture this coming Friday will be canceled> is 0.6, and his B_A for this proposition equals 0.3, then his degree of doxastic

[9] Let max(x, y) = n (where n equals the larger value of x and y, if there is such a value; if not, n=x).
[10] Thus, for ease of calculation, whenever $B_A \leq B_O$, $Ds = B_A$.

success (i.e., his *Ds*) for the proposition equals 0.3.[11] Similarly, if Omar's B_O for the proposition equals 0.6, and his B_A equals 0.9, then his degree of doxastic success (i.e., his *Ds*) for the proposition equals 0.3.[12]

Thus, I now employ the notion of a subject's *Ds* to reformulate *Prov2*, as follows:

> *Prov3* The degree to which a subject *S* knows some true proposition *P* is an increasing function of: (i) the value of *S*'s justificatory weight (i.e., *Jw*), and (ii) the value of *S*'s degree of doxastic success (i.e., *Ds*).

However, there is an important qualification that needs to be made with respect to (ii) in *Prov 3*. To see why, consider the distinction between: (1) a set of reasons that a subject *S* has for believing *P*, and (2) a set of reasons on which *S*'s degree of actual belief (that *P*) is based. This distinction corresponds to the distinction, commonly employed by epistemologists, between *propositional* and *doxastic* justification. Swain (1979) thus writes:

> It often happens that a person's belief that *h* is based upon a set of reasons *R* when believing that *h* on the basis of *R* is not justified for that person. This can happen even though the person has other reasons such that if the belief had been based upon those reasons, then the belief would have been justified.

In short, a subject *S* is *propositionally justified* in believing *P* just in case *S* has a set of reasons that adequately support *P*, whereas *S* is *doxastically justified* in believing *P* just in case *S* actually believes *P* and also bases this belief on a set of reasons that adequately support *P*. As Kvanvig (2003) notes, "Doxastic justification is what you get when you believe something for which you have propositional justification, and you base your belief on that which propositionally justifies it."

[11] $0.6 - |0.6 - 0.3| = 0.3$
[12] $0.6 - |0.6 - 0.9| = 0.3$

To see why this distinction is important for an account of graded knowledge, consider a scenario in which Sam watches a live news report indicating that a hurricane is moving towards South Carolina. Accordingly, Sam's B_O for the proposition <a hurricane is moving towards South Carolina> is fairly high, say, 0.85. Moreover, suppose that Sam's B_A for this proposition also equals 0.85. However, there's an important caveat. Sam's high B_A is not due to the news report he just watched. In fact, Sam thinks that all news reports are extremely unreliable. Rather, Sam's high B_A is due to the fact that he thinks there are angry gods in the Atlantic Ocean continually sending numerous hurricanes towards all parts of the East Coast and that other (more beneficent) gods intercept these hurricanes before they hit. In this situation, it seems rather untenable to think that the close proximity between Sam's B_A and B_O could contribute to Sam having a high degree of knowledge that *P* (i.e., as high as 0.85).

So, where is the problem? It seems to arise from the fact that the explanatory reasons on which Sam's B_A is based (i.e., the idea about gods in the Atlantic ocean) do not correspond to the normative reasons that Sam has in support of *P* (i.e., the live news report). Or more specifically, it arises from the fact that the explanatory reasons on which Sam's B_A is based are propositionally justified to a very low degree (if any positive degree).

Consider how this pertains to the wording of (ii) within *Prov3*—i.e., "the value of *S*'s degree of doxastic success (i.e., *Ds*) ". The scenario seems to suggest that what matters to *S*'s degree of knowledge is not merely an increasing function of the value of *S*'s degree of doxastic success (i.e., *S*'s *Ds*), but also an increasing function of the degree to which *S*'s *Ds* results from a correspondence between the explanatory reasons on which

S's B_A is based and the normative reasons that *S* has in support of *P*. Specifically, it seems that *S*'s degree of knowledge is an increasing function of both *S*'s *Ds* and the degree to which the explanatory reasons on which *S*'s B_A is based are propositionally justified. The main purpose of including this latter notion (i.e., the degree to which the explanatory reasons on which *S*'s B_A is based are propositionally justified) is that it ensures that *S*'s degree of doxastic success is not merely accidental (as is the case with Sam) but is instead creditable to *S*. This accords with Sosa's (2007) proposal that knowledge is *apt* belief: i.e., that knowledge is not merely accurate belief for which *S* is competent, but rather, it is accurate belief that is accurate *because* of *S*'s competence.

In light of this, how might (ii) from *Prov3* be revised so as to incorporate this idea? First, let a subject's J_E be the degree to which the subject is propositionally justified in the explanatory reasons on which the subject's B_A is based. Supposing that a subject's J_E equals 0.2 (i.e., supposing that the subject's degree of doxastic success results from a B_A based on explanatory reasons that are only 20% propositionally justified), it seems plausible to maintain that only 20% of the subject's *Ds* may count towards the subject's degree of knowledge. This suggests that the relevant doxastic value for determining a subject's degree of knowledge is the *product* of the subject's *Ds* and J_E. Accordingly, I call the product of a subject's *Ds* and J_E, the subject's *degree of doxastic weight* (abbreviated *Dw*), defined as follows:

$$Dw = Ds \cdot J_E$$

Or equivalently:

$$Dw = \max(B_O - |B_O - B_A|, 0) \cdot J_E$$

One might wonder why a subject's degree of doxastic weight shouldn't instead be calculated as:

$$Dw^* = \max(B_O - |B_O - B_A \cdot J_E|, 0)$$

The reason can be seen by considering a revised version of the Sam scenario. Suppose, again, that Sam hears about the hurricane from a reliable news source (thus his B_O still equals 0.85). His B_A also still equals 0.85, but this time the explanatory reasons on which his B_A is based consist partially of the evidence from the news report, but consist mostly of his ideas regarding the ocean gods. Thus, his J_E now equals, say, 0.3. Using Dw^*, this results in the following calculation:

$$Dw^* = \max(0.85 - |0.85 - 0.85 \cdot 0.3|, 0) = 0.255$$

With this in mind, further suppose that while still relying on the same evidence, Sam increases his B_A to an unduly high level (say, 1). Then, according to Dw^*, this serves to *increase* his overall doxastic weight:

$$Dw^* = \max(0.85 - |0.85 - 1 \cdot 0.3|, 0) = 0.3$$

Thus, D^* has the odd effect of suggesting that if S's B_A is based on explanatory reasons that are not highly justified, then S may do well to further increase her B_A even beyond that of her B_O.

By contrast, consider how the formula for Dw escapes this problem by calculating the two situations, respectively, as follows:

$$Dw = \max(0.85 - |0.85 - 0.85|, 0) \cdot 0.3 = 0.255$$

$$Dw = \max(0.85 - |0.85 - 1|, 0) \cdot 0.3 = 0.21$$

Given the above considerations, I now provide the final modification to the provisional account of graded knowledge:

> *Prov4* The degree to which a subject S knows some true proposition P is an increasing function of: (i) the value of S's justificatory weight (i.e., Jw), and (ii) the value of S's doxastic weight (i.e., Dw).

Finally, due to the widespread theoretical intuition in favor of the entailment thesis (i.e., that knowledge entails belief), an account of graded knowledge may do well to satisfy the following condition: If *S*'s degree of knowing *P* is greater than 0, then *S*'s B_A is also greater than 0. Call this the *weak graded entailment thesis*. I call it "weak" rather "strong" because a strong graded version of the entailment thesis would look something like: If *S*'s degree of knowledge that *P* is equal to *n*, then *S*'s B_A is also equal to *n*. I favor the weak interpretation of the graded entailment thesis because (unlike the strong interpretation) it allows for the possibility that one's degree of knowledge may surpass one's degree of actual belief, but restricts this possibility to cases in which one's degree of actual belief is greater than 0. Moreover, the weak interpretation may help explain cases in which people are more likely to ascribe knowledge to a subject than they are to ascribe belief to the subject (e.g., perhaps this may be understood in that the subject's degree of knowing *P* is higher than her degree of actually believing *P*).[13]

7. Degrees of knowledge

Drawing from the above discussion (Sections 3-6), I propose the following formula for determining *S*'s degree knowing *P*:[14]

$$K = \begin{cases} \text{avg}(Jw, Dw), & \text{if } B_A > 0 \\ 0, & \text{otherwise} \end{cases}$$

The lower half of the formula satisfies the weak graded entailment thesis (i.e., that if *S*'s degree of knowledge is greater than 0, then *S*'s B_A is greater than 0). The upper half of

---

[13] See Radford 1966; Beebe 2013; Murray, Sytsma, & Livengood 2013; Buckwalte, Rose, & Turri 2013.

[14] Let avg(*x*, *y*) = *n* (where *n* equals the arithmetical mean of *x* and *y*).

the formula satisfies *Prov4*: i.e., whenever *S*'s B_A is greater than 0, *S*'s degree of knowing *P* is the arithmetical mean of *Jw* and *Dw*, and thus, *S*'s degree of knowing *P* is an increasing function of both *Jw* and *Dw*, in accordance with *Prov4*.

To illustrate how the formula applies to a particular case, first recall how the notions of *Jw* and *Dw* are defined:[15]

$$Jw = \begin{cases} B_O \cdot (1 - (1 - JB_O) \cdot 2|0.5 - B_O|), & \text{if } B_O > 0.5 \\ B_O \cdot (1 - JB_O \cdot 2|0.5 - B_O|), & \text{otherwise} \end{cases}$$

$$Dw = \max(B_O - |B_O - B_A|, 0) \cdot J_E$$

Now, recall the scenario from Section 1 in which Charles has strong evidence (due to his careful and informed study) that there is an 85% chance of snow in Boston this evening, while Micah has moderate evidence (from his roommate) that there is an 85% chance of snow in Chicago this evening. Furthermore, while Charles B_A (which equals 0.85) accords with what is indicated by his evidence, Micah's B_A (which equals 0.99) goes beyond what is indicated by his evidence. Supposing that it will in fact snow in both Boston and Chicago this evening, it seems that Charles is in a better overall epistemic state than Micah.

---

[15] It might be argued that there are cases in which there is more than one B_O. For instance, suppose that someone tells you that she has an unfair coin, but she gives you no indication about the manner in which the coin is unfair. In such a case, one might think that you have two optimal degrees of belief, which are at equal but opposite distances from 0.5, such as 0.2 and 0.8. If this is the case, then *Jw* and *Dw* could be revised as follows (where B_O^{avg} is the arithmetical mean of all optimal degrees of belief, B_O^C is the optimal degree of belief that is *closest* to B_A, if there is one; and if B_A is equidistant from two optimal degrees of belief, then B_O^C equals either of them):

$$Jw = \begin{cases} B_O^{avg} \cdot (1 - (1 - JB_O) \cdot 2|0.5 - B_O^{avg}|), & \text{if } B_O^{avg} > 0.5 \\ B_O^{avg} \cdot (1 - JB_O \cdot 2|0.5 - B_O^{avg}|), & \text{otherwise} \end{cases}$$

$$Dw = \max(B_O^{avg} - |B_O^C - B_A|, 0) \cdot J_E$$

In Section 1, I argued that while standard binary accounts of knowledge are unequipped to tell us much about how such variables might contribute to one's overall epistemic state, I suggested that an account of graded knowledge might have more potential in doing this. Accordingly, consider how the above account of graded knowledge might be used to analyze the relevant variables as they pertain to the overall epistemic states of Charles and Micah.

We know that both Charles and Micah have a B_O of 0.85, and that Charles' B_A equals 0.85, whereas Micah's B_A equals 0.99. Thus, the above account of graded knowledge tells us that there are two more variables that we need to consider: JB_O and J_E. While it may be most instructive to compare numerous versions of the scenario so that the values of these additional variables may be widely varied, I choose just one set of values (due to lack of space) that should at least help illuminate the general mechanics of the proposed account of graded knowledge.

For Charles, suppose that $B_O = 0.85$ (as already indicated), $JB_O = 0.9$ (given the high strength of his evidence that there is an 85% chance of snow in Boston this evening), $B_A = 0.85$ (as already indicated), and $J_E = 0.7$ (i.e., the explanatory reasons on which his B_A is based are justified to a degree of 0.7).

For Micah, suppose that $B_O = 0.95$ (as indicated), $JB_O = 0.6$ (given the moderate strength of his evidence that there is an 85% chance of snow in Chicago this evening), $B_A = 0.99$ (as already indicated), and $J_E = 0.7$.

When the values for Charles are plugged into the relevant formulas, Charles' degree of knowledge is calculated as follows:

$$Jw = 0.85 \cdot (1 - (1 - 0.9) \cdot 2|0.5 - 0.85|) = 0.7905$$

$$Dw = \max(0.85 - |0.85 - 0.85|, 0) \cdot 0.7 = 0.595$$

$$K = \text{avg}(0.7905, 0.595) = 0.69275$$

Similarly, when the values for Micah are plugged into the relevant formulas, Micah's degree of knowledge is calculated as follows:

$$Jw = 0.85 \cdot (1 - (1 - 0.6) \cdot 2|0.5 - 0.85|) = 0.612$$

$$Dw = \max(0.85 - |0.85 - 0.99|, 0) \cdot 0.7 = 0.497$$

$$K = \text{avg}(0.612, 0.497) = 0.5545$$

In summary, given the postulated values for the relevant variables, Charles' degree of knowledge (0.69275) regarding the relevant proposition is higher than Micah's degree of knowledge (0.5545) regarding the relevant proposition. This accords with the above suggestion that Charles' overall epistemic state is better than that of Micah. Moreover, it is evident that this difference is due largely to Micah's relatively low JB_O and unduly high B_A; or, in terms of Sosa's AAA account of knowledge, the difference is due largely to Micah's relatively low competence and relatively low accuracy with respect to believing the proposition <it will snow in Chicago this evening>.

8. Conclusion

In conclusion, I began by providing two motivations for an account of graded knowledge (Section 1). I then examined various notions pertaining to belief and justification, and I considered their potential roles in an account of graded knowledge (Sections 2-6). Finally, I used the various notions of belief and justification to create a formula for determining a subject's degree of knowledge for a given proposition (Section 7).

References

Beebe, J. (2013). A Knobe effect for belief ascriptions. *The Review of Philosophy and Psychology*, 4: 235-258.

BonJour, L. (2002/2009). *Epistemology: Classic problems and contemporary responses*. Lanham MD: Rowman and Littlefield.

Buckwalter, W., Rose, D., & Turri, J. (2013). Belief through thick and thin. *Noûs,* 47 (3).

Hetherington, S. (2001). *Good knowledge, bad knowledge: On two dogmas of epistemology*. Oxford: Oxford University Press.

Hetherington, S. (2006). Knowledge's boundary problem. *Synthese*, 150: 41-56.

Hintikka, J. (2006). Epistemology without knowledge and without belief. In J. Hintikka (Ed.), *Socratic epistemology: Explorations of knowledge seeking by questioning* (pp. 11-37). Cambridge: Cambridge University Press.

Murray, D., Sytsma, J., & Livengood, J. (2013). God knows (but does God believe?). *Philosophical Studies*, 166: 83-107.

Quine, W.V. (1987). *Quiddities: An intermittently philosophical dictionary*. Cambridge MA: Harvard University Press.

Radford, C. (1966). Knowledge—by examples. *Analysis*, 27: 1–11.

Sosa, E. (1980). The raft and the pyramid: Coherence versus foundations in the theory of knowledge. *Midwest Studies in Philosophy*, 5: 3–25.

Sosa E. (1985). Knowledge and intellectual virtue. *The Monist,* 68: 226-245.

Sosa E. (1988). Methodology and apt belief. *Synthese,* 74: 415 - 426.

Sosa, E. (1991). Intellectual virtue in perspective. In E. Sosa (Ed.), *Knowledge in perspective*. Cambridge: Cambridge University Press.

Sosa, E. (2007). *Apt belief and reflective knowledge, volume 1: A virtue epistemology*. Oxford: Oxford University Press.

Sosa, E. (2009). *Apt belief and reflective knowledge, volume II.* Oxford: Oxford University Press.

Sosa, E. (2011). *Knowing full well*. Princeton, NJ: Princeton University Press.

Weiler, G. (1965). Degrees of knowledge. *The Philosophical Quarterly,* 15: 317-327.