Synthese (2012) 187:377–392
DOI 10.1007/s11229-010-9831-2

Knowledge and conviction

David James Anderson

Received: 15 April 2010 / Accepted: 9 October 2010 / Published online: 23 October 2010

Abstract Much philosophical effort has been exerted over problems having to do with the correct analysis and application of the concept of epistemic justification. While I do not wish to dispute the central place of this problem in contemporary epistemology, it seems to me that there is a general neglect of the belief condition for knowledge. In this paper I offer an analysis of 'degrees of belief' in terms of a quality I label 'conviction', go on to argue that one requires more conviction in a proposition in order to know it than to merely believe it, and conclude by suggesting that some current epistemological issues admit of new insight when we begin taking conviction seriously.

Keywords Knowledge · Belief · Degrees of belief · Conviction

1 Introduction

The claim that one must hold to a proposition more firmly in order to know it than to merely believe it is implicit in much of our thinking about the concept of knowledge, but it is rarely acknowledged and made explicit. I aim to show that it is worth thinking about carefully. This paper proceeds in three sections: In Sect. 1 I explain how I will be thinking of the doxastic attitudes, and I offer an account of the various degrees those attitudes can come in by making reference to the strength with which a proposition seems true or false (a feature I label "conviction"). In Sect. 5 I argue that knowing a proposition requires not only that one believe that proposition, but that one's degree of belief be sufficiently strong. In the final section I point toward some of the implications of this argument: (a.) it affects the important concept of epistemic defeat; (b.) it

D. J. Anderson (✉)
Purdue University, West Lafayette, IN, USA
e-mail: djanders@purdue.edu

suggests a way of defending invariantism about knowledge from the recent criticism of Contextualism; and (c.) it encourages a re-evaluation of the Closure Principle and all of the epistemic puzzles associated with it.

2 Propositional attitudes, doxastic attitudes, and degree of strength

I will use the term 'proposition' to refer to the primary bearer of truth-value. I assume that propositions are abstract entities which are expressed by certain of our declarative sentences, but if the primary bearers of truth value are sentence tokens or something else, then relevant minor alterations should render all of my comments applicable. To consider a proposition is simply to hold it before one's mind and understand what it claims. Consideration is distinct from mere understanding since one can understand a proposition at some time t even if one is thinking entirely of other things at t. Consideration requires the direction of attention in addition to understanding. I take the concept of *understanding a proposition* to be basic and unanalyzable. I also take it to be a necessary condition for genuinely considering a proposition. One might consider the words of a sentence without understanding which proposition those words express. One might also understand a sentence to be expressing one proposition when in fact it could be understood to be expressing another (this may be what happens when one has a 'minimal' understanding of something deep and complex but doesn't appreciate all of the subtleties involved). One cannot, I submit, consider a proposition itself unless one understands it. Consideration is a combination of understanding and directed attention.

We can consider propositions for a variety of different purposes. An historian of philosophy might consider the proposition that *there exists a material world* for the sole purpose of determining which early modern thinkers accepted it. One who studies ancient linguistics might consider the proposition *the cat is on the mat* for the purpose of determining which of several possible Sanskrit sentences best expresses it. Most importantly for my purposes, we often consider propositions for the sake of determining their truth value.

When we consider a proposition for the purpose of determining its truth value (but not when we consider it for some other purpose), our consideration of it is often attended with a certain hard-to-define mental state; a 'seeming' of a very precise sort. For some propositions (e.g. *I have hands*) we describe this mental state by saying that it "seems true". For others (e.g. *snow is purple*), we describe it by saying that it "seems false". Following William Tolhurst[1] I will say that a proposition seems true when its consideration is attended with a "felt veridicality" that recommends or demands a response of affirmation. "The notion of felt veridicality," claims Tolhurst, "resists analysis. But we all are acquainted with it; reflection on suitable examples should be sufficient to call it to mind."[2] Felt veridicality is simply the feeling of agreeableness

[1] Tolhurst (1998). Note that Tolhurst focuses on perceptual seemings while I want to focus on propositional seemings more generally, but the analysis seems to work in either case. I borrow Tolhurst's term "felt veridicality" and part of his analysis of "seeming true", but what follows are my own refinements to the view.

[2] *Ibid.*, p. 299.

which is shared in common by any proposition about which you would say that it "seems true." Tolhurst says nothing about negative seemings, but in a similar way I will say that a proposition seems false when its consideration is attended with a "felt falsidicality" (for lack of a better antonym of "veridicality"[3]) which recommends or demands a response of denial.

Felt veridicality comes in degrees, and because of this one proposition can "seem true" to a greater degree than does another. Both $2 + 3 = 5$ and *Barack Obama is the current president of the United States* seem true to me, and I recognize a felt demand to affirm both, but the felt veridicality of the former is greater than that of the latter. It can even be the case that, among two propositions that one believes to be necessary, one can seem true in a stronger sense. $2 + 3 = 5$ seems true to me in a stronger way than does *there are infinitely many primes*. As truth itself does not come in degrees, it is *not* the case that I have a full strength seeming that "$2 + 3 = 5$ is true to the degree n" and a full-strength seeming that "*I have hands* is true to degree $n - m$". Rather, my seeming that "*I have hands* is (fully) true" is less powerful than is my seeming that "$2 + 3 = 5$ is (fully) true".

Also because felt veridicality comes in degrees, one's consideration of some propositions can be attended with a weak felt veridicality that merely suggests, but does not demand, a response of affirmation. As an example, consider the proposition that *President Obama will be re-elected*. If you are anything like me you would hesitate to say that this proposition "seems true"; rather, you might describe the mental state attending your consideration of the proposition by saying that it "seems more likely to be true than to be false". We can say corresponding things for felt falsidicality: propositions whose consideration is attended with a very strong felt falsidicality will seem false to a greater degree than will those with a somewhat lesser degree of felt falsidicality. There are also propositions which "seem more likely to be false than to be true"; these are ones whose consideration is attended with a (relatively) weak felt falsidicality which merely *suggests* denial rather than demands it.

It is important to distinguish prima facie felt veridicality from *all things considered* felt veridicality.[4] Consider the Monty Hall problem:

> Suppose you're on a game show, and you're given the choice of three doors: Behind one door is a car; behind the others, goats. You pick a door, say No. 1, and the host, who knows what's behind the doors, opens another door, say No. 3, which has a goat. He then says to you, "Do you want to pick door No. 2?" Is it to your advantage to switch your choice?[5]

For most people, an initial consideration of the proposition *it is to the contestant's advantage to switch his choice* is attended with a felt falsidicality, perhaps strongly enough that it demands a response of denial. After reflecting a little more on the situation and calculating the odds, however, one comes to realize that the contestant

[3] The Oxford Dictionary of Philosophy defines "falsidical" as "an uncommon opposite of veridical" (p. 130).

[4] Likewise for felt falsidicality.

[5] Whitaker (1990).

actually has a 2/3 chance of winning if he switches doors. At this point, if one reconsiders the italicized proposition, he may still recognize *some* felt falsidicality. In some sense, it still 'seems' that it shouldn't matter which door the contestant goes with. This felt falsidicality, however, is outweighed by a strong felt veridicality coming from the calculations just completed. I will say that, although we have some prima facie felt falsidicality toward the proposition, our *all things considered* mental state is one of strong felt veridicality.

In some cases one's all things considered response ends up being balanced out by equally strong felt falsidicality and felt veridicality. An example: suppose your friend is accused of murder. As you look into the case you realize that there is some very strong evidence for his guilt. At the same time, he is a close friend with a gentle disposition and you have a very hard time reconciling what you know about his character with the details of the crime. When you consider the proposition *my friend committed murder*, it may be that a prima facie felt veridicality stemming from your recognition of the prosecution's evidence is balanced out by a prima facie felt falsidicality stemming from your knowledge of your friend's character. *All things considered* the proposition is attended with neither a felt veridicality nor a felt falsidicality. In other cases one's all things considered response is neutral not because the prima facie seemings balance each other out, but because upon consideration the proposition has absolutely no felt veridicality or felt falsidicality. When a flipped coin lands out of sight one's consideration of the proposition *the coin landed on heads* is attended with neither a felt veridicality nor a felt falsidicality; and neither affirmation nor denial is even slightly recommended.

It will be useful to have a single term to refer to these types of propositional seemings. In what follows I will speak of one's "conviction" in or against a proposition. Conviction is always one's all things considered response to considering a proposition for the purpose of determining truth value. Conviction can be positive (when consideration is attended with all things considered felt veridicality), negative (when attended with all things considered felt falsidicality), or neutral (when attended with neither, either because, like the coin flip case, there just are no seemings either way, or because, like the murder case, the prima facie positive seemings are balanced out by the prima facie negative ones). Positive conviction and negative conviction are mental states that come in degrees. I will speak of positive conviction *in* a proposition and negative conviction *against* a proposition. So, one has strong positive conviction in p when one's consideration of p is attended with a particular strong felt veridicality; one has weak positive conviction in p when one's consideration of p is characterized by a weaker felt veridicality. Correspondingly, one can have strong or weak negative conviction against a proposition, when one's consideration of the proposition is attended with (strong or weak) felt falsidicality. One has neutral conviction in a proposition when one considers it for the purpose of determining truth value, but, all things considered, one has neither any felt veridicality nor felt falsidicality toward it. Neutral conviction is thus not a unique mental state over and above consideration, and it does not come in degrees. Considering a proposition is thus necessary but not sufficient for having some type (and degree) of conviction in it. It is necessary because propositions that are not being considered do not have any sort of seeming associated with them. It is not sufficient because it is possible to conceptually entertain a proposition for some purpose other than determining whether or not it is true.

Although I define conviction as a very particular sort of 'seeming', the former term is preferable because of the latter's ambiguity. Seemings can be prima facie or *all things considered* but conviction is always one's total, all things considered response. When one becomes convinced that some proposition p entails a contradiction, for example, one's negative conviction in p is very high (maximally high in ideal cases) regardless of any prima facie seemings of veridicality. Also, while many seemings are based in perception, there are also seemings that are more broadly cognitive. When one considers a straight stick partially submerged in water one typically experiences both a perceptual seeming as of a bent stick, and a cognitive seeming that the stick is not really bent. Conviction is simply the total felt veridicality or felt falsidicality of a proposition, regardless of whether it is based in perception or other cognitive mechanisms. Finally, conviction can be neutral (when one has considered a proposition in the relevant manner but has neither positive nor negative conviction) but one does not have neutral seemings.

In addition to merely conceiving of, or cognitively grasping, a proposition (with or without some degree of conviction), we take various attitudes toward propositions. Consider the proposition that *it will be 95 degrees this afternoon*. As I prefer a much cooler climate, I might take an attitude of fear toward that proposition. My wife, who prefers warmer climates, might instead take an attitude of hope toward the same proposition. Hope and fear, then, are attitudes that we can take to propositions. We can also take an attitude of affirmation or denial toward a proposition. The former attitude is what we call belief and the latter is disbelief. The different attitudes one might take to a proposition are "propositional attitudes." Belief, disbelief, hope, and fear are all propositional attitudes. Belief and disbelief, however, are members of an important subset of the propositional attitudes which are concerned with the truth or falsity of the proposition. Members of this subset are the "doxastic attitudes".

Doxastic attitudes (as well as attitudes more generally) are sometimes explicitly adopted and consciously attended to, but they are present implicitly even when they are not being attended to. A victim of arachnophobia, for example, has a fear of spiders even when she is wholly occupied with other matters in an environment with no spiders around. Likewise, in addition to believing right now that Barack Obama is the president of the United States (because I am currently considering that proposition), I also believed it an hour ago when I was thinking only about getting a cup of coffee. A distinction is often drawn between what one believes *occurrently* and what one believes *dispositionally*. The former are the few beliefs that, at any one time, one is consciously attending to. The latter compose the host of other things that we believe even when not thinking about them. One has conviction only in one's occurrent beliefs, since propositions one is not attending to do not have any sort of felt veridicality or falsidicality. Nevertheless, it seems to me that all of one's dispositional beliefs were, at one time, held occurrently. What I want to say about the relations between conviction and the doxastic attitudes can extend to dispositional beliefs by considering the time at which they were occurrent, and the strength of the believer's conviction at that time.

However, because I've limited dispositional beliefs to those that were once occurrent, we must make finer distinctions about our beliefs. This is because there are many propositions we have never consciously considered (and so have never believed occurrently), and yet which we typically claim to believe. To borrow an example from

Robert Audi,[6] consider the proposition *98.165 is larger than 98*. If I were to ask you right now if you believe that proposition you would unhesitatingly respond in the affirmative. Furthermore, if I were to ask you whether you believed it a few moments before reading this paragraph, you would likely also respond in the affirmative. We can suppose, however, that before reading this paragraph you have never considered that particular proposition, and so have never before believed it occurrently. Under that assumption, I've committed myself to the claim that you didn't believe it dispositionally either.I claim that, before reading this paragraph, you did not believe that 98.165 is larger than 98. Have I just insulted your intelligence?

Before you take offense, consider the distinction between 'having a dispositional belief that *p*', and 'being disposed to believe that *p*'. In learning mathematics you formed actual beliefs about many principles such as *adding a fractional part to any number yields a larger number*. Those actual beliefs (which were dispositional rather than occurrent just before you read the previous paragraph) disposed you to form the belief that *98.165 is larger than 98* immediately upon considering it. What occurs is that, owing to certain of your actual dispositional beliefs, upon your consideration of the proposition you are immediately struck with a strong felt veridicality that causes the belief. It is not as if you needed to consult the actual beliefs you did hold and then draw an inference to your freshly acquired belief that *98.165 is larger than 98*. Rather, your actual beliefs made it the case that this proposition strongly seemed true as soon as you considered it. So, one is disposed to believe *p* when one does not actually believe it but (because of various facts about oneself, including things that one does actually believe) one would come to believe *p* (rather than disbelieve or withhold judgment about it) in suitable circumstances. At any given time, one occurrently believes no more than a handful of propositions, one dispositionally believes very many other propositions which are 'stored' in memory after having once been occurrently believed, and one is disposed to believe very many more things under suitable circumstances.[7,8]

---

6 Audi (1994, p. 419).

7 Because it would take me too far afield, I omit any discussion of what it takes for circumstances to be "suitable." It will not include all possible circumstances. An infant is not disposed to believe that *98.165 is larger than 98* despite the fact that, were the infant granted full conceptual powers by God and a range of beliefs about the relevant mathematical principles, the infant would believe that *98.165 is larger than 98*. For more discussion on this topic, I refer the reader to Audi (1994).

8 One objection to the account just given (suggested to me by an anonymous referee) is that our behavior demonstrates that we *actually* believe propositions we have never considered (and hence that do *not* count as actual beliefs on my account). As an example, I often read a book as I walk from my office to the parking lot. During this walk my attention is almost entirely on the book, and yet I am able to navigate around fire hydrants and other people walking on the sidewalk. My ability to do so, it might be thought, is evidence that I *do* believe things like (a) *the fire hydrant is about 5 feet in front of me* and (b) *if I continue walking along this path, I will run into the lady jogging towards me*. However, those are propositions that I never consciously consider, toward which I have no conviction at all, and which, consequently, I never occurrently believe. If all actual beliefs are either occurrent or dispositional, and all dispositional beliefs are occurrently held at some time in the past, then I do not actually believe (a) and (b) at all. But if that is the case, how is it that I am able to *use* them in guiding my behavior?

I am willing to bite the bullet of this objection and maintain that I do not form any belief in (a) or (b). It does not seem to me a difficult bullet to bite. I think of belief as a cognitive attitude we take toward propositions. We use them to make rational judgments and in guiding our actions. Much of our behavior,

Just like many other propositional attitudes (including hope and fear), the doxastic attitudes each come in degrees of strength. I hope both that *my children grow to be strong and healthy* and that *we don't have leftovers for dinner again*, but I hope much more strongly for the former. Similarly, I disbelieve both that $2 + 3 = 4$ and that *mereological universalism is true*, but I disbelieve the former much more strongly. Accounting for these degrees can be a very difficult endeavor.[9] It is difficult not just for the case of belief but for the case of any other type of attitude. What does it mean to say that I hope for one thing more strongly than I hope for another, or that I like apple pie more than I like peach pie? Answers to the latter questions will, I think, go a long way toward providing an answer to the question about various strengths of belief. Counterfactuals can often be a useful tool for identifying stronger or weaker beliefs (i.e. a stronger belief is one that would be retained in the face of more or stronger counter-evidence than would a weaker one[10]), but given the notorious difficulty of defending counterfactual analyses of anything else in epistemology, I am skeptical that they can be used to *analyze* strength of belief.

Many philosophers attempt to analyze degrees of belief as credence functions that (in a perfectly rational agent) obey the probability calculus. According to this analysis, when a perfectly rational agent throws a normal six-sided die and considers the probable outcome, he does not *fully* believe that *the die will not land on a 6*. His degree of belief in that proposition rather takes into account the probabilities involved and he believes the proposition to a degree of 5/6, or 0.83. A rational agent thus proportions his degree of assent to the strength of the supporting evidence.[11]

I am skeptical about the need for distinguishing what we might call 'partial' belief from what we might call 'full' or 'flat-out' belief. Certainly we do (or at least *ought to*) treat the proposition *the die will not land on a 6* differently than we treat the proposition *the die will not land on a 5 or a 6*, but I do not think that our attitude toward either of those propositions counts as the affirmation characteristic of belief. The correct thing to say about my relation toward those two propositions seems not to be that "I believe the first to the degree 0.83 and the second only to the degree 0.66", but rather that "I believe neither of them". I think this contention is supported by introspection on our doxastic attitudes. There is a difference in type, and not merely in degree, between

Footnote 8 continued
however, does not count as intentional acting and does not need to be informed only by our beliefs. In denying that I have any beliefs about the locations of fire hydrants and people as I walk to my car, I am not denying that I have any sort of access to that information. I am merely claiming that the information is processed at a subconscious level and never attains the status of a belief. I think it is important to distinguish the information about the world that we acquire and store as attitudes towards propositions, from the information that we use in mundane, every-day tasks but do not develop explicit attitudes toward. I am using the term 'belief' to refer exclusively to the former class, and I do not take this to be a non-standard use of the term. The latter class seems more important for our procedural knowledge (or 'knowledge-how') than our propositional knowledge (or 'knowledge-that'). My ability to ride a bicycle implies some sort of access to information about the necessary leg and arm movements for riding a bike, but I do not think it implies that we have any *beliefs* about that information.

9 In 2009, *Synthese* published a book-length collection of essays devoted to this very topic.

10 See Gärdenfors and Makinson (1988) on 'epistemic entrenchment'.

11 There are too many supporters of this view to try to list them all, but for a prominent account see Swinburne (2001).

the attitude we take towards propositions we accept, and the attitude we take toward propositions that we merely consider likely. For example, I have a vague memory of rain during my last birthday and I am thus favorable toward the proposition that *it was raining on my last birthday*. Because my memory is vague enough that it may be the previous birthday that I am remembering, however, I do not accept it the way I do the proposition *it was raining yesterday* (of which I am quite certain). The attitude I bear to propositions I accept is the one we most properly refer to with 'belief,' even though we sometimes find it useful to speak loosely about 'believing' propositions we are merely favorable towards. I do not accept either of the propositions about the outcome of a die roll in the way that I accept that it was raining yesterday; and the difference is not merely one of degree.

If I believe neither *the die will not land on a 6* nor *the die will not land on a 5 or a 6*, however, how are we to explain the difference in my behavior regarding the two propositions (e.g. that I would accept different bets on whether or not they are true)? It seems to me that we can adequately explain this difference in at least two ways: (a.) My behavior might simply be a result of my attitude of (full) belief in the propositions *the chance of the die landing on anything but a six is 0.83* and *the chance of the die landing on anything but a five or a six is 0.66*;[12] (b.) My behavior might not be explained with reference to belief at all, but rather with reference to my degree of conviction in the two propositions. Although I believe neither proposition about the outcome of the die roll, I would accept different bets on them because I have greater conviction in the one than in the other.

One believes a proposition p more strongly than one believes a different proposition q when (i.) one has (fairly) strong positive conviction in both p and q; and (ii.) one has greater positive conviction in p than in q.[13] Note that while one might have a particular degree of conviction in p and also an explicit belief about the probability that p is true, these two mental states are not identical. After all, one might have a particular degree of conviction in p without having any belief at all about how likely p is. One's conviction in p is not itself a belief (it is a different sort of mental state). One's conviction in p is simply a measure of how strong one's felt veridicality or felt falsidicality in p is.

We can account for the difference in the way that we act with regards to distinct propositions, none of which we believe, by considering the relative strength of our conviction in those propositions. We can also, I think, account for our stronger belief in some propositions and our weaker belief in others with the same consideration. Finally, I think that focusing on conviction gives us a way to account for the difference between propositions we fully affirm (believe), and those we think somewhat likely but do not believe. Both positive and negative conviction range on a scale from a maximal to a minimal seeming (of veridicality or falsidicality, respectively). In between minimal negative conviction and minimal positive conviction lie some propositions for which we have neutral conviction. There is some threshold on the lower end of negative conviction below which we form an attitude of rejection and disbelieve the proposition;

[12] Plantinga (1993, p. 8) endorses this response.

[13] For an account of degrees of belief very close to mine that uses the notion of 'confidence', see Foley (1992).

and there is also some threshold on the upper end of positive conviction above which we form an attitude of affirmation and believe the proposition. We neither believe nor disbelieve a proposition when our conviction in it lies between these two thresholds.

The thresholds in the degree of conviction may be vague, and may also be relative to the believer and/or to different contexts. Because the thresholds are relative to person and context, a particularly gullible person will affirm a proposition with considerably less conviction than will a particularly skeptical person. If they have vague boundaries, it may be that our conviction in some propositions is so close to the threshold that we lack introspective access to whether or not we believe it. It may be that our concepts of belief and disbelief, like our concepts of baldness and smoothness, simply lack the refinement needed to provide a definite answer for all cases. I do not take this vagueness to be especially problematic, nor do I think it a reason for rejecting the notion of conviction thresholds for belief and disbelief. The concepts are precise enough for our purposes.

In addition to affirming and denying a proposition, one can withhold judgment altogether. It is thus not uncommon to distinguish from both belief and disbelief a third member of the set of doxastic attitudes.[14] When a proposition's felt veridicality is strong enough, we take the favorable, affirming attitude of belief toward it. When, on the other hand, its felt falsidicality is strong enough we take the unfavorable, denying attitude of disbelief toward it. For many propositions, however, our degree of conviction in them is not strong enough either to demand a response of affirmation or a response of denial. We withhold judgment about such propositions.

One can withhold judgment in different ways. One might withhold on two distinct propositions, but at the same time be much closer to affirming one than the other. For example, when I consider the proposition *there are over one million blades of grass on my lawn* I find that I am very close to affirming it. However, I've never counted the blades of grass in a representative portion of my lawn nor multiplied that result by the proportion of the lawn it represents. I consider the proposition very likely, but I think it would be a mistake to claim that I take an attitude of affirmation toward it. I am much closer to affirming it than denying it, but I nevertheless fail to have either attitude. In such cases I will say that one 'weakly withholds' the proposition. On the other hand, I have no inclination whatsoever to believe that *there is an even number of blades of grass on my lawn*, nor have I any inclination to disbelieve it. I have neutral conviction in that proposition. I will say that one withholds a proposition more strongly the closer he is to having neutral conviction in the proposition. So, I weakly withhold that *there are over a million blades of grass on my lawn* and I strongly withhold that *there are an even number of blades of grass on my lawn*.

Since the strength of conviction can be used to analyze degrees of belief/disbelief, and also the various distinctions between ways that one can fail either to believe or to disbelieve, it is an important notion to be clear on. One important question is: in virtue of what is one's conviction in a proposition as strong as it is? How is strength of conviction determined? It seems that many factors can affect conviction. In ideal cases (for many of our beliefs at least) the most (only?) significant factor will be one's

[14] See, for example, Bergmann (2005, p. 421).

evidence or reasons for that proposition. My conviction in the proposition that *Jones is the murderer* increases as I come to know that his fingerprints are all over the murder weapon, that the victim was recently caught in bed with Jones' wife, and that Jones has no credible alibi. It would increase even more if I witness video footage of Jones committing the act of murder. However, unfortunate though it may be, we are not ideal rational agents. One's felt veridicality or falsidicality toward a proposition can vary not only with different evidence, but with different desires, volitions, and general moods. One's conviction in something like the Principle of Sufficient Reason, for example, is surely affected by arguments one has considered which either support or discredit it; but it can also be affected just by the amount of time one has recently spent reading Leibniz, Spinoza, and other rationalists. Had one instead spent that time reading early modern empiricists it may well be that one's felt veridicality when considering PSR would simply not be as strong. Again, when one desires a certain outcome very much, one's conviction in propositions suggesting that an alternate outcome will instead be realized might be much weaker than it would otherwise be.[15]

Although there are many ways that the strength of our conviction in a proposition can be affected, our doxastic attitude toward a proposition is determined by the degree of our conviction. Sufficiently strong conviction in some proposition p causes an attitude of belief in p. When one gathers further evidence resulting in a weaker conviction in p, one naturally abandons one's belief and instead withholds judgment. If even further evidence results in strong negative conviction, one comes to disbelieve p. This change in attitude happens automatically, below the level of consciousness. Forming a belief or disbelief in a proposition is thus not an action on our part, but rather a response to our conviction in it. The relation here is a causal one. Having high conviction in p is conceptually distinct from believing p, so it is not logically necessary that one believes whenever one's conviction is sufficiently high. Nevertheless, I think the two concepts are connected tightly enough that it would be very odd[16] to claim that a proposition seems very strongly to be true to a certain person *and* that he does not believe the proposition.[17]

[15] It is because conviction can be affected by various factors that it has a *normative* dimension. When considering some belief, it is important to ask not only how strong one's conviction in the believed proposition is, but also how strong it *ought* to be. Accounts of epistemic justification typically focus merely on the conditions under which someone is justified in believing a proposition. What I've said suggests that these accounts need to be extended to give conditions under which one is justified in having some particular degree of conviction in a proposition.

[16] Perhaps one could cognitively malfunction in such a way that one forms beliefs in response to relatively weak positive conviction. Perhaps it is even possible for one to malfunction to the extent that one forms beliefs even in response to negative conviction. Furthermore, perhaps it is possible for there to be rational beings whose cognition functions in such a way that they form beliefs without conviction at all. When I talk of a certain degree of conviction *causing* certain doxastic attitudes, the discussion is implicitly restricted to the scope of healthy, properly functioning human beings.

[17] My position entails doxastic involuntarism—the thesis that we do not have voluntary control over our own doxastic attitudes—but I think that many of the voluntarist's more powerful criticisms are answered by my account of conviction and the factors affecting it. While one does not have direct voluntary control over the type of attitude one forms in response to conviction, one does have some voluntary control over the strength of conviction in the first place. If I am enamored with a particular position, I can refuse to honestly consider arguments against it, and thus retain my high conviction in it.

The picture I am drawing, then, is one in which our consideration of certain propositions is attended with a felt veridicality or a felt falsidicality that comes with more or less strength. When the felt falsidicality is very strong the proposition seems false, and we respond with an attitude of rejection toward it (this is disbelief). I then posit some threshold of conviction at which there is still a felt falsidicality toward the proposition, but which is not so strong that we would say it "seems false." Although we remain closer to denying these propositions than to affirming them, we form neither attitude and instead withhold judgment. As the felt falsidicality toward a proposition becomes less and less strong (as one's negative conviction in it grows weaker and weaker), we become less and less inclined to deny it. There is a point at which its consideration is attended with neither felt falsidicality nor felt veridicality, and at this point I say that one has neutral conviction and strongly withholds judgment on the proposition. On the positive side, one becomes and more and more inclined to affirm a proposition as its consideration is attended with stronger and stronger felt veridicality. Again, I posit some threshold of conviction at which our inclination to affirm overcomes our withholding of judgment and we form an attitude of belief. Degree of belief, then, is a matter of how much one's conviction exceeds that threshold for belief. With this picture of the doxastic attitudes and their various strengths in mind, we are prepared to consider the degree of belief required for knowledge.

3 Knowledge and degree of belief

It is common to analyze our concept of knowledge into four components:
A subject *S* knows a proposition *p* iff:

(1) *S* believes *p*
(2) *p* is true
(3) *S* is justified in believing *p*
(4) Some anti-Gettier condition is satisfied.

Much debate is centered on the proper formulation of the condition described in (4). Much debate is also centered on how to best understand the nature of the justification required for knowledge in (3). I hope to avoid these debates altogether and instead suggest that, even if conditions (3) and (4) are satisfactorily formulated, the above analysis of knowledge is incomplete. More specifically, I want to argue that condition (1), as stated, is too weak. Not only must one believe *p* in order to know *p*, but one must believe it with (at least) some minimal degree of strength. On at least some readings of condition (3) it is possible, or so I want to maintain, that one meet all of conditions (1)–(4) and yet fail to have knowledge of *p* because one's belief in *p* is too weak.

The argument goes like this: Prior to analysis, there is a strong intuition that it is constitutive of knowledge that it have some stability for the knower. One does not reject his belief in a known proposition on the basis of evidence or reasons he considers to be relatively weak; rather, when a proposition is known, one's attitude toward that proposition does not change without reasons or evidence that the knower takes to be relatively strong. This stability is one of the reasons that we value knowledge more

than we do mere true belief. Now, it is commonly assumed that the presence of doxastic justification[18] (or of strong enough doxastic justification to warrant a belief) will sufficiently account for the stability required for knowledge. It is often claimed that one has sufficient doxastic justification for knowing p when one has sufficient propositional justification for knowing p, and one's belief that p is also based on whatever provides that propositional justification.[19] As such, cases in which one has sufficiently strong doxastic justification for knowing p coupled with a very weak belief in p are rare. Nevertheless, it is worth pointing out that there is nothing in principle keeping these situations from being possible. Suppose, then, that some subject S believes a proposition p, p is true, S is strongly justified in believing p, and S is not in a Gettier situation. Suppose further that (perhaps because she under-values the strength of her own evidence, or perhaps simply because she doesn't want p to be true) S's conviction in p is only minimally strong enough to result in her believing it. p seems true to her with enough strength that she takes an attitude of affirmation toward it, but she remains very close to lacking that affirmation, and hence she nearly withholds judgment on p instead. In such a case, it seems to me that if knowledge requires no more conviction in the known proposition than does belief, then we shall have to say that S knows p. This is very difficult to accept once we appreciate the instability of S's belief in p. Having only the minimal degree of conviction necessary for belief, S will lose her attitude toward p with only minimal reason to doubt p. We shall have to say that, although S knows p, S can lose that knowledge very easily. The degree to which we find this result unacceptable is the degree to which we should accept that knowledge requires a higher degree of conviction than does belief.

Consider the following intuition pump: Jones has taken out a second mortgage to place a bet on his favorite racehorse Papa Clem. Unfortunately for him, nature calls just as Papa Clem's race begins and Jones spends the duration of the race in the restroom. Upon returning to the stands he gathers some very strong evidence supporting the proposition that

(A) *Papa Clem lost the race*.

Perhaps he overhears some people complaining about the money they just lost on Papa Clem, sees a portion of the race recap on the television in which Papa Clem is lagging far behind on the final lap, and notices a group of people consoling Papa Clem's owner. Jones adopts an attitude of affirmation toward (A), but does so with reluctance; I stipulate that this attitude is very weak. The proposition seems true to Jones, but his confidence in Papa Clem's superiority and his fear at losing all of his money keep his felt veridicality from being as strong as his evidence would demand. We may say that Jones has the minimal degree of conviction in (A) for having an attitude of belief in it. As a result, Jones' belief in (A) is very unstable. If Jones were to receive even some very weak reason for thinking that (A) is false (some non-committal testimony from a questionable witness, for example), he would lose his belief in (A) and would withhold judgment on it instead. Now, despite the weakness of Jones' belief in (A),

[18] See Bergmann (2006, p. 4) for the distinction between propositional and doxastic justification.

[19] See, for example, Conee and Feldman (1985).

he does believe it. We can stipulate that (A) is also true, it seems that Jones has strong doxastic justification for (A),[20] and we can also stipulate that Jones is not in a Gettier situation. Does Jones *know* (A)?

My own intuition about this example is that he does not have knowledge. Jones' lack of knowledge is not a result of his having insufficient justification, nor is it a result of his belief not being grounded properly on that which justifies him. There is no objectionable luck keeping Jones' belief from amounting to knowledge, nor is he subject to any other Gettier-type conditions. The only thing keeping Jones from knowing that Papa Clem lost, I submit, is that although his degree of conviction in the proposition is sufficiently high that his attitude toward it counts as belief, it is not sufficiently high to meet the demands of knowledge. Knowledge, I conclude, requires a higher degree of conviction than does mere belief.

Note that I am not arguing that knowledge requires maximal conviction, or certainty. I think that I know both that *modus ponens is a valid inference form* and that *I ate an orange for breakfast last Tuesday*; nevertheless my conviction in the former proposition is considerably higher than is my conviction in the latter. Lower still is my conviction that *mereological universalism is false*, but while I believe this last proposition I do not claim to know it. I claim that knowledge requires a high degree of conviction but maintain that one can have more conviction in one known proposition than in another.

I've already posited a threshold of conviction below which we form the doxastic attitude of disbelief and above which we withhold judgment, and also a threshold of conviction below which we withhold and above which we believe. I now posit a third threshold, not between the attitude of belief and some distinct attitude we might call 'super-belief', but between attitudes of belief which are not eligible candidates for knowledge and attitudes of belief which are. Only true, justified, non-Gettiered beliefs which meet or exceed this threshold amount to knowledge.

In light of this discussion, we might reformulate the conditions for knowledge like this:

S knows that *p* iff:

(1*) *p* is true.
(2*) *S* believes *p*, and *S*'s conviction in *p* meets or exceeds the threshold for knowledge.
(3*) *S* is justified in having sufficiently strong conviction in *p* to meet the threshold for knowledge.
(4*) Some anti-Gettier condition is satisfied.

I have several comments about this formulation. First, conditions (2*) and (3*) make reference to the conviction threshold for knowledge. This threshold is just the strength that one's conviction must have in order for the belief produced by it to be

[20] If it does not seem that his evidence is strong enough to ground knowledge, we can modify the example by making it stronger. Doing so will make it less likely that Jones' affectations would keep his conviction at the minimal level for belief, but there is no reason to think it impossible.

eligible for knowledge. Since knowledge requires more stability than mere belief, this threshold will be above that level of conviction which produces belief. Were we able to quantify conviction and pinpoint the precise degree required for knowledge, then we would be able to refer to this threshold without using the term "knowledge". Because of this, the analysis is not circular.

Second, (2*) requires that S "meet or exceed" the conviction threshold for knowledge, but condition (3*) only requires that S be justified in having the minimal degree of conviction required for knowledge. Jointly, these conditions allow for S to believe a proposition more strongly than her evidence justifies, and for that belief still to count as knowledge. Suppose that Peter's evidence for some proposition p is strong enough to justify a conviction that meets the minimal requirements for knowledge of p, but also that (because of non-rational factors) Peter's conviction in p is maximal (that he believes p with certainty). My analysis delivers the result that Peter knows p, and I think that this is the result we want (even though his overconfidence is irrational).

Finally, note that the anti-Gettier condition we add in (4*) must be strong enough to rule out not only the standard sort of Gettier cases, but also cases like the following: suppose that Peter meets conditions (1*) and (3*) for some proposition p. Then, p is true and Peter is justified in having sufficiently strong conviction to know that p. Suppose furthermore that, because of non-rational factors, Peter's conviction at some time t_1 is lower than it ought to be and he fails to believe p. At t_1, then, Peter fails to meet condition (2*) and so fails to know that p. But now suppose that, through other non-rational factors, Peter's conviction at a later time t_2 is raised to a level sufficient for knowledge. He thus comes to believe that p and to meet condition (2*). Peter's belief is a result of a double case of non-rational factors affecting his conviction, and thus seems too accidental (or at least non-rational) to count as knowledge. An appropriate filling out of condition (4*) should cover such concerns.

4 Implications

My thesis that knowledge requires more conviction than does mere belief has some interesting ramifications for several issues in contemporary epistemology. Lacking the space to work out these ramifications here, I will merely gesture toward them with the suggestion that future research would be profitable.

First, if a relatively high degree of conviction is necessary for knowledge, then there are conditions under which one loses knowledge that are not covered by familiar accounts of epistemological defeat. Most discussion of defeat, or defeaters, focuses on the question of how and when a justified belief loses its status as justified. So long as justification is conceived of as a necessary component of knowledge, part of the interest of a theory of defeat is its application to the project of understanding the conditions under which we lose knowledge (i.e. every justification-defeater is, ipso facto, a defeater of knowledge as well). If it is true that knowledge requires a higher degree of conviction than does mere belief, however, then the project of understanding the conditions under which we lose knowledge will be incomplete if we focus only on justification defeaters. There will be cases, in other words, in which our knowledge, or warrant, is defeated without our justification being defeated. In addition to losing

knowledge when one's justification is rebutted or undercut,[21] one will lose knowledge when (for any reason you like) one's conviction is lowered below the threshold required for knowledge. A full theory of epistemic defeat will have to account for and explain these "conviction defeaters".

Second, I've said that the threshold of conviction above which belief is produced is (to a certain degree) person-relative. A generally gullible person might form beliefs with less conviction than does a generally skeptical person. I've also said that the level of conviction required for knowledge might vary with context; so that a particular degree of conviction might be sufficient for knowledge in one context but not in another. If this is right, it suggests a way of defending a classical invariantism about the standards of justification required for knowledge. Classical invariantists claim that we have a single definite conception of knowledge, and that the standards of justification required for a belief to amount to knowledge are fixed. Contextualists disagree, claiming that "knowledge" is an indexical term which denotes different concepts in different conversational contexts. A statement such as "Jones knows where his car is parked" might express a true proposition when made in a conversational context with relatively relaxed standards, and might also express a falsehood when asserted in certain philosophical conversations.[22] Others respond to classical invariantism by claiming that, while the term "knowledge" univocally refers to a single concept, the standards of justification required for a belief to amount to knowledge can vary depending on the believer's situation.[23] Both of these rejections of classical invariantism are typically motivated by examples and thought experiments in which we have competing intuitions about whether or not the subject has knowledge.

If what I have said about conviction is correct, then perhaps we can propose a satisfying analysis of the examples used to support these denials of invariantism while maintaining both that "knowledge" refers to a single concept *and* that the justificational standards required for knowledge do not change. "High stakes" situations in which it seems wrong to claim that the subject has knowledge might be analyzed as situations in which one needs *very* strong conviction in order to know, rather than as situations in which one needs very strong justification in order to know, or in which the term "knows" picks out a very stringent concept of knowledge.

Finally, if knowledge requires a certain degree of conviction, lack of conviction in certain cases might require a revision of the Closure Principle. Roughly, Closure says that, whenever a person (i.) knows some p; (ii.) knows that p entails q; and (iii.) forms a belief in q on the basis of (i) and (ii), that person also knows that q. Closure enjoys a lot of intuitive support because it seems clear that known logical implication preserves justification; if I have enough justification in p for belief in it to count as knowledge, and I also *know* that p entails q, then if I base a belief in q on $[p \ \& \ (p \rightarrow q)]$, it seems that I should have just as strong justification in q. If knowledge requires more conviction than does belief, however, we can endorse these strong intuitions and still deny this rough formulation of Closure. It may happen that I have full knowledge

21 See Pollock (1986, p. 196) for the classic distinction between these two types of defeaters.

22 See DeRose (1995) for a defense of contextualism.

23 See, for example, Stanley (2005).

both of p and $(p \rightarrow q)$, and that I derive a belief in q on that basis, but that (for whatever non-rational reason you like) I simply do not have strong enough conviction in q to meet the requirements for knowledge. Since the Closure Principle is related to many puzzles and problems in epistemology (skepticism and lottery problems to name two), new light might be shed on these puzzles once we get clear about how conviction affects the closure of knowledge under known entailment.

5 Conclusion

I have argued that believing a proposition p requires that one's consideration of p be attended with a strong enough felt veridicality. I have argued further that knowing p requires this felt veridicality to be even stronger. Finally, I have pointed to three ramifications of these claims and suggested ways that they may shed light on some of the puzzles and problems that face contemporary epistemology. Our understanding of knowledge and its related concepts can be improved, I suggest, by further exploring the notion of conviction.

References

Audi, R. (1994). Dispositional beliefs and dispositions to believe. *Nous, 28*(4), 419–434.
Bergmann, M. (2005). Defeaters and higher level requirements. *The Philosophical Quarterly, 55*, 419–436.
Bergmann, M. (2006). *Justification without awareness* (272 pp.). New York: OUP.
Blackburn, S. (2005). *The Oxford dictionary of philosophy* (407 pp.). New York: OUP.
Conee, E., & Feldman, R. (1985). Evidentialism. *Philosophical Studies, 48*, 15–34.
DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review, 104*(1), 1–52.
Foley, R. (1992). The epistemology of belief and the epistemology of degrees of belief. *American Philosophical Quarterly, 29*(2), 111–124.
Gärdenfors, P., & Makinson, D. (1988). Revisions of knowledge systems using epistemic entrenchment. In M. Vardi (Ed.), *Theoretical aspects of reasoning about knowledge* (pp. 83–95). Los Altos: Morgan Kaufmann.
Plantinga, A. (1993). *Warrant and proper function* (243 pp.). New York: OUP.
Pollock, J. (1986). *Contemporary theories of knowledge*. Savage: Rowman and Littlefield Publishers.
Stanley, J. (2005). *Knowledge and practical interests*. New York: OUP.
Swinburne, R. (2001). *Epistemic justification* (272 pp.). New York: OUP.
Tolhurst, W. (1998). Seemings. *American Philosophical Quarterly, 35*(3), 293–302.
Whitaker, C. F. (9 September 1990). "Ask Marilyn" column. *Parade Magazine*, p. 16.