Brent Emerson
15 December 1997

Contextualism in Action
k-knowledge and the Myth of Bivalent Epistemic Status

1. Introduction

In the words of epistemologist David Lewis, "We know a lot."[1] It's clear there is also 'a lot' that we don't know. Epistemologists (and people in general) seem to like to talk about this distinction in a *bivalent* sense: that is, either we know something, or we do not know it. Only two discrete epistemic alternatives are available. This is perhaps strongly supported by a tradition that has viewed knowledge as 'true justified belief'—we generally speak as if an assertion is either true or not, justified or not, believed or not. The bivalence of knowledge would seem to simply follow. Other epistemic terms are treated this way as well: in a contextualist consideration of what counts as knowledge and what constitutes a challenge to it, Keith DeRose has treated various bivalent concepts—*relevance*, *sensitivity*, etc.—all either applying to statements absolutely or not at all.[2] The bivalence of epistemic status seems to be universally taken for granted.

The problem is that this bivalent view is untenable. The important concept of *how well* we know something makes clear that knowledge is located on a continuum, not in a binary switch. Bivalent epistemic status is a myth that ultimately serves only to confuse us and make our use of the concept of 'knowledge' seem incoherent. Appealing to the details of knowledge-acquisition processes as they actually occur, this paper will introduce multivalent epistemic status (*k*-knowledge) by sketching and roughly formalizing mathematically a general, contextualist theory of its use. It will also examine how a theory

---

[1] Lewis, David. 'Elusive Knowledge' *Australasian Journal of Philosophy*, 74 (December 1996): 549. All Lewis citations are from this article.
[2] DeRose, Keith. 'Solving the Skeptical Problem' *The Philosophical Review*, 104 (January 1995): 1–52. All DeRose citations are from this article.

of multivalent knowledge can explain the place of epistemological skepticism and formalize a refutation that relates to strategies both contextualist and naturalist.

2. Actual Knowledge of the Actual World

It's easy to see that two discrete predications of the concept 'know' are not enough. It's certainly important to us to determine whether we know things, but the qualifying factor of *how well* we know them is just as important. Without more than two 'knowledge-values', we wouldn't be able to ask questions like 'how well do you know (that *A*)?' But we do (and need to) ask questions like this all the time.[3] Even some epistemologists seem to mention this multivalence, if vaguely and without realizing its implications. Lewis, for instance—at the beginning of a paper about presumably bivalent knowledge—speaks of "those things that we know *better than* we know the premises of any philosophical argument to the contrary."(my emphasis)[4] DeRose uses the "notion *of (relative) strength of epistemic position*"[5] without paying formal attention to what this means. If this multivalence is meant to be implied (in these papers or in epistemological investigation in general), it might at least be argued that it isn't given the proper attention.

To begin, then, we assume that a main goal of epistemology is to effectively model actual knowledge-acquisition processes, to explain what knowledge is and how we come by it. This view is clearly empiricist rather than normative, though any normative theory should be so grounded in the actual if it is to be relevant (i.e., have actual application.)

We assume further that epistemic investigation is always undertaken for the purpose of best satisfying an epistemic goal. And epistemic subjects have different goals. The goal

---

[3] A possible objection: questions like "how well do you know?", in requiring the kind of multivalent answer they require, actually ask us to determine the probability of our bivalent-knowledge of *A*. Then knowledge is bivalent after all. It's not that we don't have perfect epistemic access to the external world, we just don't have perfect internal access to whether we know something or not. Response: (*a*) this is not a substantive difference and (*b*) it is wrong, since the multivalence in the situation actually is, in practice, encapsulated in our concept of knowledge. Saying otherwise simply means you're using 'knowledge' differently than most people use it.

[4] Lewis, p.549.

may be access to the 'truth', prediction of future observation, prediction of what everyone else will agree with, or something else. The details are for now unimportant: this goal will be denoted by 'obtains'. Something will be said to **obtain** if it satisfies the epistemic subject *S*'s goals. Assuming that we want to maintain a bivalent notion of what obtains, we can still have a theory of knowledge that admits of degrees, as long as we deny perfect epistemic access to this external bivalent world[6]. The knowledge we are left with is *k*-knowledge. As a definition:

S ***k*-knows** that *A* if and only if *S* estimates that *k* is the probability that *A* obtains.[7]

This entails our first rule, the

Rule of Multivalence
All knowledge admits of degree.

What remains is the question: how do we come to know things better or worse than we know other things? How do we gain *k*-knowledge? To get an answer we need a reasonably formal way of talking about it.

3. Formalizing *k*-knowledge

We can think of the knowledge-acquisition process as an attempt to make a subject's knowledge-bank most correctly correspond to what obtains—an attempt to hold knowledge of maximal *k*. To accomplish this, an epistemic subject *S* needs a decision procedure for how well she knows a certain assertion *A*. Toward this end, we locate the variability in how-well-*S*-knows-that-*A* on [-1,1], a portion of the real number line:

---

[5] DeRose, p.29.
[6] A non-bivalent epistemology like this could probably be used to problematize further the notion of a bivalent ontology, but that is not the present project.
[7] At this point it seems that *k*-knowledge is relativized to the subject, which it is. This does not necessarily entail relativity to the individual, as a social group of individuals could perfectly well constitute an epistemic subject. Knowledge can be held by a group of people of any size to the extent that they share the common epistemic goal of finding out what obtains.

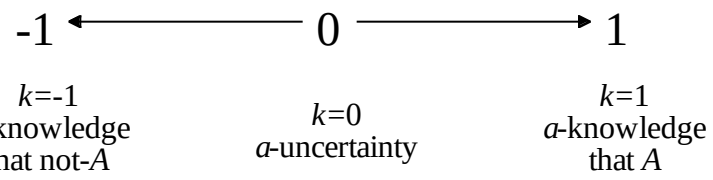

Here *a(bsolute)*-knowledge of *A* is mapped to 1, complete ignorance (*a*-uncertainty) of *A* to 0, and *a*-knowledge of *A*'s negation to -1. What will interest us, of course, are the values of *k* that fall between these extremes.

We are now in a position to consider how it is that *S* acquires *k*-knowledge. If we make the assumption/restriction that the only way for *S* to acquire knowledge is by consulting authorities, we might imagine her in the world entering into various **knowledge-acquisition processes**, consulting various authorities made available to her by her sensory faculties. This would entail the use of processes in which she places a great deal of trust (linguistically-mediated information from professors at major universities, visual information from her own eyes) as well as those which she trusts mildly (information from the commercial news media, reports from strangers) and others in which she places negligible trust (idle speculation from her parents from when she was young and inquisitive, reports from people she thinks are stupid). She might trust a math professor to give better information (than a high school student) about mathematical questions and the local high school student's reports to lead to better knowledge (than the math professor) about which bands were coming to town soon. Not surprisingly, it makes sense that process-authority (how well *S* trusts these processes) is multivalent in exactly the same way as knowledge: *S* authoritizes particular knowledge-acquisition processes (over certain fields of knowledge) to the degree that *S* trusts them to give her knowledge (within the relevant fields)[8]. Each

[8] It could be argued that *S* could always misplace her trust, leading to a flaw in this formulation—if knowledge-acquisition is to be formalized like this, process-authority should be determined not by subjective trust, but by a process's tendency to reliably represent what actually obtains. To make this argument, however, would be to engage in serious question-begging. The point of introducing *p*-authority is so that *S* has a procedure for acquiring useful knowledge and justifying it to others. Given that nobody has perfect epistemic access to what obtains, an 'objective' system of p-authorities would deny *S* access to the values of *p* needed to determine to what extent her knowledge is valuable. Indeed, the purpose of the

process *P* thus attains a ***p*-authority**, where *p* corresponds to the degree *k* of the knowledge *S* would obtain if *P* were the only process involved. The variability in how-well-*S*-trusts-process-*P* is located on [0,1]:

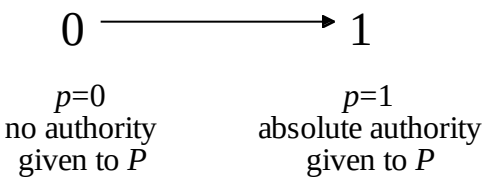

A process *S* trusts completely is given an authority of 1, a process she totally distrusts an authority of 0.

But to gain useful knowledge usually requires of *S* that she rely on several different knowledge-acquisition processes. To get *k*-knowledge of some assertion *A* from a book in a library, for instance, S has to enter into many processes—understanding the book, trusting that the authors actually are tenured professors at M.I.T. and Princeton, believing that well-published people at high-profile universities usually know what they're talking about, etc.—each of which she assigns an independent process authority *p*. It is clear that the knowledge she gets from the **acquisition-route** ('reading this particular book, in these particular present circumstances') that includes these processes should be trusted only to the extent that the processes themselves are. Thus the

Rule of (Multiplicative) Composition
To calculate $k_i(A)$—the degree of knowledge of *A* obtained via a particular i^{th} acquisition-route—multiply together all the authorities for different processes involved in that route:

$$k_i(A) = \prod_j p_j$$

One further complication: in her attempt to determine what obtains, *S* usually consults several independent acquisition-routes: asking a smart friend, checking several books, listening to a professor in lecture, thinking about it on her own, etc. We may

---

project in the first place is to determine what obtains; if the details of 'what obtains' are made availably prior to epistemic investigation, such investigation becomes pointless.

assume, then, that she arrives (via the *Rule of Composition*) at several different answers to the question 'To what degree do I know that *A*?' How is she to combine them into a single best estimate that *A* obtains? Here we assume[9] that *S* combines her best reason for believing that *A* with her best reason for not-*A*, via the

Rule of (Additive) Combination
To calculate *k(A)*—the ultimate degree *k* to which an assertion *A* is known, in the presence of more than one acquisition-route—add together the maximum positive degree k_i^+ and the minimum negative degree k_i^-:

$$k(A) = \max\left(k_i^+(A)\right) + \min\left(k_i^-(A)\right)$$

4. An Example

One morning, *S* sees Keith walking toward her in the post office. "Why didn't you call me back?" he says. *S* doesn't know what he's talking about, and she tells him so. "I called you this weekend," he says, "why didn't you call me back?" Now, *S* is very sure that Keith didn't call her this weekend—she would have talked to him or received a message from one of her housemates or from a message pad or answering machine. Still, *S* is not completely (*a*-) sure that he didn't call, just (say) 80%-sure. (Her housemates are *k*-known to sometimes not give messages, etc.) *S* wishes that she *k*-knew enough about Keith to *p*-authoritize him on this matter, but she just met him, and as a rule she disregards information from individuals or processes of indeterminate *p*-authority. 80%-knowledge is in this case enough for *S* to make a fairly certain 'obtains-determination', and she is about to sneer "I

[9] Other rules of combination could be proposed; I flirted myself with one that combines *all* the answers *S* gets rather than just the extremal ones. I find this extremal one the most plausible, however, given that I find plausible examples like this: imagine that *S* has information from her (uneducated) friend that 'France is located in Europe'. When she consults a geography book in the library, *S* is looking for an answer she can better trust—one that more precisely approximates *a*-knowledge. Once she has the more reliable information, her friend's help no longer matters. Justifying her knowledge that 'France is located in Europe', *S* invokes only her book-process and not her only-somewhat-trustworthy-friend-process. Only the extremally trusted acquisition-route enters *S*'s judgment and justification of her knowledge. Alternately, we can make the difference between the two rules trivial: in the case where two people who are trusted equally give the same answer, *S* might trust the process of 'asking these two people of *p*-authority' to give her better than *p*-knowledge because there is more than one person involved. Then *S* has a new acquisition-route which she trusts better than the routes that consist of asking the two people individually.

know you didn't call me," and walk away. Luckily, Keith stops her by pulling a letter from his pocket. It reads:

> To Whom It May Concern,
>
> Keith is one of my good friends. He is intelligent and a fast runner and, if he says that he called you at a certain time, he indeed did call you then.
>
> Sincerely,
> Senali

Ah, this changes things: the letter provides a new route whereby *S* can get knowledge from Keith. It appears that Senali, one of *S*'s good friends, *p*-authoritizes Keith to *k*-know about whether or not he has called someone at a certain time, to tell the truth about it, etc. *S* can't be *a*-sure that Senali actually wrote the letter (as she can't be *a*-sure of *anything*), but she *k*-knows that Senali did (it features her signature and letterhead, which are reasonably hard to forge). Knowing Senali as *S* does, she *p*-authoritizes her on matters of people knowing whether or not they called other people.

Now *S* has a decision/calculation to make. Via her 1st knowledge-acquisition route, *S* already 80%-knows that Keith didn't call her. Via a 2nd knowledge-acquisition route, *S* can use information from Senali to *p*-authoritize Keith and get *k*-knowledge from him. If *C* is Keith's calling of *S*, *S* *x*-knows that Senali wrote the letter, estimates that Senali *y*-authoritizes Keith, and herself *z*-authoritizes Senali on matters like this, *S* can (using the *Rule of Composition*) make the following inference:

> 1. I *x*-know that Senali *y*-authoritizes Keith over a field of knowledge that includes *C*.
> 2. I *z*-authoritize Senali over the field of knowledge that includes Keith's knowledge of *C*.
>
> _____
> I *xyz*-know that *C*.

This process thus gives *xyz*-knowledge that *C*. The last step involves a judgment of how well *S* trusts this process: supposing that she *u*-authoritizes this way of getting information, *S* can conclude that she *uxyz*-knows *C*. Say *uxyz*=74%. Then, using the *Rule of Combination*, in the absence of any other knowledge-acquisition processes, *S* (80% - 74%

=) 6%-knows that Keith didn't call her this weekend. In this case, it turns out that *S* trusts her friends just a little more than this elaborate play with the letter and all; she makes some polite excuse and rushes off to more important things. If it had turned out that *S k*-knew better that Keith was lying (if this ultimate *k* had been larger), she might have simply turned her back on Keith and walked away, or been even more curt: we might suppose that her rudeness would increase proportionally with her *k*-knowledge that Keith was lying to her.[10] And *S* makes judgments like this one every day!

5. Contextualism in Action

A confusion in the above example might arise at the point of S's calculation of the degree of her knowledge via the second (complicated) route. The confusion: why stop at *uxyz*-knowledge? Shouldn't there be a factor for how well *S* trusts the entire process, and then one for how well *that* larger process is trusted, and on and on, *ad infinitum*? Further, shouldn't *S* calculate with additional internal factors for how well she knows that the words in the letter mean what she thinks they mean, for how well she knows that the words she read were the words that were actually written (how well she trusts the quality of lighting in the room, her own eyesight, the care she took in reading), again *ad infinitum*?

This is where the contextualist, armed with the concept of *k*-knowledge, steps in. In specifying a context of knowledge-evaluation (or -acquisition), says the contextualist, we pretend to find infallible a series of such processes by giving them absolute authority. A **context** consists in the setting of an (infinite) series of *p*'s to 1 (i.e., of absolute trust placed in a series of knowledge-gathering processes), leaving us a finite series of processes to consider. This definition entails our:

[10] The *k*-formalism has been meticulously observed for good reason, but is now getting tiresome, I'm sure. From now on the *k*- prefix will be omitted except when necessary. All epistemic terms used should

Rule of Context

To calculate *k(A)*, set authorities p_j to 1 for the knowledge-acquisition processes which define the context in which *A* is currently being evaluated (i.e., ignore the possible fallibility of these processes by pretending they are infallible).

$$p_{n\cdots\infty} = 1$$

With this formalism, we can analyze the skeptical attack and the contextualist's response. The skeptic (says the contextualist) raises the standards of knowledge-attribution to levels that are unreachably high, thereby precluding any possibility for knowledge. But a complete theory should also explain how this raising of the standards is accomplished[11], which we are now in a position to do. All the skeptic has to do to raise these standards is to present the hypothesis that no knowledge-acquisition process is to be trusted completely. Then, using the *Rule of Composition*, every *k* is equal to a multiplication of an infinite series of authorities between 0 and 1; hence every *k* equals 0. Knowledge is completely impossible.

The contextualist responds simply: in presenting this skeptical hypothesis, the skeptic is denying the *Rule of Context*. The skeptic is trying to speak from an acontextual position—the skeptical conclusion is actually that *acontextual* knowledge is impossible. Since context is required for knowledge, acontextual knowledge is certainly not possible, as it is incoherent. Contextual knowledge is therefore safe from the skeptic.

Lemma: A Brief History of *S*'s Knowledge that *H*

> For no sooner do we engage in epistemology—the systematic philosophical examination of knowledge—than we meet a compelling argument that we know next to nothing. The skeptical argument is nothing new or fancy. It is just this: it seems as if knowledge must be by definition infallible. If you claim that *S* knows that *P*, and yet you grant that *S* cannot eliminate a certain possibility in which not-*P*, it certainly seems as if you have granted that *S* does not after all know that *P*. To speak of fallible knowledge, of knowledge despite uneliminated possibilities of error, just *sounds* contradictory.[12]

---

nonetheless be understood to contain it.

[11] This seems to be the contextualists' main challenge, referenced in DeRose, Lewis, Schiffer.

[12] Lewis, p.549.

The above passage is a perfect illustration of how skepticism looks in a bivalent epistemic paradigm, where "it seems as if knowledge must be by definition infallible." We can identify two classes of bivalent subject: the infallible subject (who challenges skeptical arguments and believes that knowledge is possible) and the skeptical subject (who claims that "*S* cannot eliminate a certain possibility in which not-*P*"). With no room for fallible knowledge (which is nonsensical in a bivalent paradigm), skepticism (i.e., the suggestion of the lack of absolute knowledge) entails the impossibility of knowledge. Here's how skepticism looks in a *multivalent* paradigm:

epistemic subject	S_i	EPIPHANY OF ERROR	S_f	EPIPHANY OF SKEPTICISM	S_s
subject's assumed status of knowledge	infallible		fallible		impossible
process hypothesis	$\forall p : p = 1$		$\exists p : p \neq 1$		$\forall p : p \neq 1$
epistemic hypothesis	$\forall k : \pm k = 1, 0$		$\exists k : 0 < \pm k < 1$		$\forall k : k = 0$

It is fair to assume that *S* begins life without a sense that her knowledge of the world could be incorrect. As S_i, *S* naively believes her knowledge to be infallible; she places absolute trust (p=1) in all of her knowledge-acquisition processes. With no reason to doubt this 'ones-all-the-way-down' hypothesis, information comes for her in only two varieties: known (k=±1) or unknown (k=0). As S_i, her sensory information *is* her world: she knows that 'here is a hand' (*H*) because here *is* a hand, and vice versa. She sees the ontological bivalence of the external world as coextensive with her epistemic bivalence.

At some point, however, S_i will invariably become S_f. When *S* is introduced to the notion of error, it becomes clear that her knowledge-gathering faculties sometimes lead her to err about the status of the external world. Ontological bivalence suddenly becomes prior to epistemic status—her epistemic investigation can only be viewed as an approximation of the absolute external. As such, her knowledge of ordinary everyday statements like *H*

seems fallible: since knowledge-acquisition processes are no longer to be trusted absolutely, knowledge occurs only in degrees. If skepticism is to be defined (as bivalence would have it) as 'non-absolute knowledge of *H'*, S_f is skeptical—we might identify her as the 'mild skeptic', since she does succeed in acquiring contextual *k*-knowledge. In this sense every epistemic subject (who is not naive in the sense of S_i nor completely skeptical in the sense of S_s) is a 'skeptic'. We might even say that this mild 'skepticism' is *required* for epistemic investigation; what is skepticism in a bivalent paradigm passes as reasonable fallible knowledge given multivalence.

The final blow to knowledge occurs when S_f encounters skeptical arguments to the effect that *no* knowledge-acquisition process should be trusted absolutely. If S_f is thus convinced to search for acontextual knowledge (with no p=1), she will find that all her knowledge is completely uncertain (k=0). As S_s, she is completely skeptical. This is the skepticism that multivalent contextualism, by insisting on the *Rule of Context*, refutes. It is clear to the contextualist that knowledge is *actually* fallible in the way that S_f considers it to be—if a normative conclusion is warranted, this is the epistemic position in which *S* should find herself.

6. The Naturalist Connection

> [Our] theory of the world may work well in enabling us to predict and deal with experience, but what reason have we to accept that it is *true*? Quine's response, again, is to raise the issue of the standpoint from which the critic speaks. The critic's standpoint is not, evidently, that of our ordinary system of the world, since the truth of the whole of that theory is being cast in doubt. But neither, in asking the question, does Quine's critic have in mind another system of the world which can underpin her use of the notion of truth: in that case we would simply have a conflict of theories, to be settled on the ordinary sort of scientific grounds. So Quine's critic seems to speak neither from within our current theory of the world, nor from within an alternative. But for Quine the point of naturalism is that we are always within a system of the world, and that our talk of truth is tied to that system.[13]

To make the connection to naturalized epistemology (à la Quine), we can simply identify the critic's *standpoint* with her *context*. Naturalism tells us that the skeptical subject—one who

seems to entertain global doubts—has no standpoint to speak from. The local doubts that naturalism expects science (and hence life in general) to require can be identified with the (finite number of) p_j that S_f does not set equal to 1. Local doubts (such as S_f maintains) make for a context in which knowledge can be acquired, for a standpoint from which the rest of the epistemic landscape can be investigated and questioned. Global doubts (S_s's insistence that no p_j are equal to 1) make for an inherently impossible attempt to speak and acquire knowledge acontextually, (i.e., without a standpoint).

S_i	S_f	S_s
no doubts; non-'skeptical'	local doubts; mildly 'skeptical'	global doubts; completely 'skeptical'

Again, the only reasonable place to be is at S_f.

7. The Bivalent/Multivalent Incoherence

The skeptical question might be said to reduce to this: *how does* S *'know' anything?* The difficulty appears when *S* maintains that she knows 'a lot' while simultaneously being compelled by skeptical arguments to the contrary. The difficulty disappears when we realize that *S* is operating (out in the world) with a continuum-based criterion of knowledge while theorizing on a bivalent basis. Our concept of knowledge is sure to thus become incoherent if *S* deals ordinarily in *k*-knowledge but occasionally adopts the assumptions of the bivalent paradigm and seeks *a*-knowledge. This suggests an important question: how does *S* confuse *k*-knowledge for *a*-knowledge? Why ever step into the bivalent paradigm?

Without getting overly psychological, we might speculate that the short period of time *S* spends as S_i colors her epistemic assumptions for her entire life. *S* might admit in practice the defeasibility of her perceptual beliefs while still maintaining more deeply the strongly impressed belief in a perfect epistemic access to the absolute. Or perhaps our social conventions confuse us: in using a concept something like 'good-enough-for-this-

[13] Hylton, P. 'Quine's Naturalism' *Midwest Studies in Philosophy* XIX (1994):264.

situation' and agreeing on a *k* which splits the knowledge-continuum into two distinct portions, *S* might produce a sort of contextual pseudo-bivalence. (One can imagine a group of scientists instituting standards for what they'll accept as experimentally verified, reducing a multivalent continuum of 'how-well-verified' to the bivalent status of 'published or not-published'. In a case like this, *S* might confuse 'above-*k*' knowledge with *a*-knowledge, 'below-*k*' with *a*-uncertainty. Clearly these distinctions are central; a better-informed investigation into the origins of binarization of epistemic status would be an important contribution to the type of epistemological theory we've discussed.

In summary, we can say that the presented theory is **multivalent** in that an assertion *A* can be known in degree (contrast with the 'absolutely' of epistemic bivalence). It is **contextualist** in that *A* can be known (whether in the acquiring or the attribution) in different contexts (depending on which acquisition-processes we choose to pretend to trust absolutely). It is **naturalist** in that it denies the existence of *a priori* knowledge (all knowledge is necessarily gained via experience of the world) and epistemic investigation is made into a science of determining how best to satisfy a subject's epistemic goals. A mild, 'healthy' skepticism is actually that which enables our continued epistemic investigation; an extreme form untenable and nonsensical. Knowledge, it is concluded, is incoherent only when as epistemic subjects we confuse our multivalent knowledge for a (mythical) bivalent type, seeking acontextual knowledge when context is integral to knowledge from the beginning.