

Knowledge and Approximate Knowledge

Lieven Decock

Faculty of Philosophy, VU University Amsterdam
l.b.decock@vu.nl

Igor Douven

Faculty of Philosophy, University of Groningen
i.e.j.douven@rug.nl

Christoph Kelp

Institute of Philosophy, University of Leuven
christoph.kelp@hiw.kuleuven.be

Sylvia Wenmackers

Faculty of Philosophy, University of Groningen
s.wenmackers@rug.nl

Abstract

Traditionally, epistemologists have held that only truth-related factors matter in the question of whether a subject can be said to know a proposition. Various philosophers have recently departed from this doctrine by claiming that the answer to this question also depends on practical concerns. They take this move to be warranted by the fact that people's knowledge attributions appear sensitive to contextual variation, in particular variation due to differing stakes. This paper proposes an alternative explanation of the aforementioned fact, one that allows us to stick to the orthodoxy. The alternative applies the conceptual spaces approach to the concept of knowledge. With knowledge conceived of spatially, the variability in knowledge attributions follows from recent work on identity, according to which our standards for judging things (including concepts) to be identical are context-dependent. On the proposal to be made, it depends on what is at stake in a context whether it is worth distinguishing between knowing and being at least close to knowing.

For most of the history of epistemology, its practitioners have held that only truth-related factors matter to knowledge. Various philosophers have recently departed from this doctrine by claiming that whether or not we know a proposition depends at least in part on other factors as well, most notably, on what the practical consequences would be for us were we mistaken about the proposition. This would imply that what we know may vary from one context to another, given that what is at stake for us in being

right about a proposition may depend on the practical environment we are in. There is a wealth of evidence to be found in the literature—including, most famously, our pre-theoretical responses to DeRose’s [1992:913] Bank Cases and Cohen’s [1999:58] Airport Cases—that supports the context-dependence of knowledge by showing that people’s willingness to attribute knowledge to a given subject is sensitive to contextual variation, in particular contextual variation due to differing practical concerns. (For ease of reference, from here on we will call this evidence “the variability data.”)

The support is *indirect* inasmuch as the evidence in itself only indicates that our *attributions* of knowledge are sensitive to contextual variation. Perhaps this is best explained by assuming that knowledge is somehow dependent on practical concerns, but it cannot be excluded off-hand that better explanations of the variability data are to be had.

Given that in this paper we aim to advance an alternative explanation, it is good to now make a general comment on how to proceed when there is more than one hypothesis that accommodates the data. In the philosophy of science, this question has been addressed under the heading of underdetermination. It is generally (even if not universally) accepted that if two hypotheses stand in the same logical relationship to the available evidence—for instance, if both entail the evidence, or (weaker) if they are both consistent with the evidence (but neither entails the evidence)—then choosing between them is a matter of which strikes the best balance in terms of satisfying the so-called theoretical virtues, like simplicity, fruitfulness, elegance, and coherence with background theories (see Douven [2008a]).

We side with the mainstream on this point. Our goal is to present an explanation of the dependence of our knowledge attributions on pragmatic factors that does better on precisely such theoretical grounds than the alternatives already on offer. In particular, we argue that the data at issue are best explained by drawing on an account of concept identification that flows from work proposed on independent grounds in Douven and Decock [2010]. The explanation, while also in pragmatic terms, will be seen to leave intact the received view of knowledge as being conceptually detached from practical matters and, more generally, as being context-insensitive. Moreover, this explanation does not single out knowledge as being somehow special; the relevant considerations of concept identification pertain equally to other concepts.

We begin by briefly discussing the main extant explanations of the data at issue here, pointing out what we think are their main shortcomings. The account of identity on which our own alternative explanation is based assumes the conceptual spaces approach, a geometric view of conceptualization that is increasingly popular in the cognitive sciences. To be able to apply this account of identity to the concept of knowledge, we must show that knowledge can be represented geometrically, and adequately so. We present the conceptual spaces approach in Section 2, and in that section we also show how knowledge can be modelled in spatial terms. After describing the account of identity that we will use (Section 3), we apply that account to the concept of knowledge, thought of *more geometrico*, and we show how it follows from this application that knowledge attributions are sensitive to what is at stake in a context, without any

implication that knowledge itself is pragmatically encroached or otherwise sensitive to context (Section 4). Roughly, the central claim of the paper is going to be that it depends on what is at stake in a given context whether it is worth distinguishing in that context between a person's knowing a proposition and that person's being at least close to knowing the proposition.

1. Contextualism, pragmatic encroachment, or loose talk? Cohen [1987], [1988], [1999], DeRose [1992], [1995], and Lewis [1996] were among the first to take seriously the variability data. They proposed as an explanation for these data a view that is now commonly known as "contextualism." In this view, knowledge attributions may vary with context because in different contexts we are not, or at least need not be, attributing the same thing. The extension of "know" and related words may vary from one context to the next so that knowledge-attributing sentences may express different propositions in different contexts. To these authors' minds, this not only accounts for our intuitive responses to the Bank Cases and similar scenarios, it also explains why in everyday contexts we attribute to ourselves and others knowledge of a host of mundane facts while we tend to retract most of these attributions when confronted by the skeptic.

In the meantime, various problems with this position have been reported in the literature, the main ones certainly being that contextualism's proposed solution to the skeptical problem remains unsatisfactory (e.g., Klein [2000], Kornblith [2000], Sosa [2000], Feldman [2001], Oakley [2001]); that contextualism is committed to the claim that competent speakers are systematically mistaken about certain semantic facts, to wit, ones concerning the standards of epistemic evaluation in contexts other than their own (e.g., Schiffer [1996], Rysiew [2001], Hawthorne [2004]); and that contextualism makes "knowledge" emerge as too much of a linguistic oddball (e.g., Douven [2004], Hawthorne [2004], Stanley [2005]). Because of these and related problems, many of its erstwhile advocates seem to have lost their enthusiasm for contextualism.

Another response elicited by the data concerning the variability of knowledge attributions, one that currently enjoys much greater popularity and that will serve as our main target in the following, is what has been called the "pragmatic encroachment view." According to pragmatic encroachers, what we know is dependent on context not because the words "know" and "knowledge" have different referents in different contexts—they do not—but because knowledge and action are conceptually interconnected. This broad idea has been spelled out in a variety of ways in the literature. To mention but the most prominent proposals in this vein, according to Hawthorne [2004], something qualifies as knowledge iff it can serve as a premise in one's practical deliberations, and Fantl and McGrath [2009] hold that one knows a proposition only if it is proper to act on the proposition. Granted that, as these authors argue, acceptability as a premise in practical deliberations and propriety to act on a proposition, respectively, are context-relative issues, the designated proposals predict the kind of context-sensitivity of knowledge attributions that we find. For present concerns, the differences in detail between Hawthorne's, Fantl and McGrath's, and similar proposals

by other authors are of little interest. What is important to note is that, as argued in Douven [2008b], all proposals of this sort are at odds with our best theory of rational agency, to wit, Bayesian decision theory.

Bayesian decision theory is not only, in Lewis's [1981:5] words, "simple, elegant, powerful, and conceptually economical," it is also at the heart of microeconomics and thereby enjoys strong empirical support.¹ According to this theory, rational agents maximize expected utilities, that is, probability-weighted averages of the utilities of the various possible outcomes of their options. Even if the nature of the relevant notion of probability as well as that of the notion of utility are subject to some debate, none of the standard explications of probability nor any of the standard explications of utility postulate an interesting connection, analytical or otherwise, with knowledge. Thus, our best theory of rational agency accords no role to knowledge. Should we really believe that there is an intrinsic connection between knowledge and agency, as Hawthorne, Fantl, McGrath, and others hold, and thus abandon an empirically successful and widely accepted theory, just because people's knowledge attributions appear to be sensitive to context? The present explanation certainly does not score high on coherence with background theories.

Naturally, this and other theoretical virtues become relevant only if there is more than one empirically adequate account of whichever the data are that one seeks to explain. In the present case, are there any serious contenders?

An alternative explanation—or rather "explaining-away"—that does not have the drawback of being at loggerheads with Bayesian decision theory argues that our knowledge attributions are only *seemingly* dependent on context. What is really context-dependent, in this account, are our standards of precision in communicating, and this makes it appear as though our knowledge attributions are context-dependent.² When the standards of precision are high, our apparent knowledge attributions are precisely that: attributions of knowledge. But often, we speak only loosely. When we do, what seem to be attributions of knowledge are really attributions of knowledge *or something close to it*. And the standards of precision vary with what is at stake in a context: if the stakes are high, it generally pays to be precise; if the stakes are low, there is much less reason to be fussy.³

This explanation is not only more conservative than either contextualism or the pragmatic encroachment explanation—for instance, it is perfectly compatible with Bayesian decision theory—it also fits nicely with a general view on language use that linguists and other theorists have been urging. As Krifka [2002] argues, for pragmatic

¹Bayesian decision theory is sometimes criticized on the grounds that it is highly idealized. But, first, the same is true for virtually any scientific theory. Second, there is a promising ongoing de-idealization program; see in particular Weirich's work in this vein (e.g., Weirich [2004]).

²See Unger [1975] and, for a more recent treatment, Davis [2007]. Lewis [1979] also observes the context-dependence of standards of precision, but for him these have semantic rather than (merely) pragmatic significance. Lewis [1996] is naturally read as developing this idea into a contextualist semantics of "knows."

³See also Kennedy [2007:24 f] on the relation between standards of precision and what is at stake in a context.

reasons, we often prefer to be brief and vague in our language use rather than precise but more longwinded (see also van Deemter [2010, Ch. 11]). Often, but not always: it is a matter of making a trade-off. It may be perfectly alright to assert

(1) The plane is scheduled to depart at 6 AM.

or

(2) It takes two hours by train from Brussels to London.

even if the plane is scheduled to depart at 6:02 AM, and even if it takes two hours and four minutes to get from Brussels to London. But sometimes that is *not* alright; neither (1) nor (2) may be an acceptable assertion in a context in which the addressee is on a tight schedule (e.g., because she has a connecting train to catch). The very same mechanism—balancing brevity and accuracy, depending on contextual stakes—may underlie the seeming variation in our willingness to attribute knowledge, depending on contextual stakes.

While the loose talking explanation has been much less in the philosophical lime-light than the two previously mentioned views, it is—we believe—much closer to the truth about our knowledge attributions. Still, not all is well for this alternative explanation. Compare these two dialogues:

Well, no

S1: It takes two hours by train from Brussels to London.

S2: Exactly two hours?

S1: Well, no, I mean approximately two hours; I think it takes two hours and a couple of minutes.

Yeah, really

S1: We could go to the bank tomorrow instead of today.

S2: Tomorrow is Saturday; do you know whether the bank is open then?

S1: Yes, it is.

S2: Do you really know that?

S1: Yeah, really!

The *Well, no* dialogue illustrates a common feature of our talking loosely in a given situation, to wit, that in such a situation typically not much pressure is needed to make us concede *that* we are speaking loosely. But this is not true of contexts in which, according to the hypothesis we are considering, we are using “know” or “knowledge” only loosely. We may imagine that, for speakers S1 and S2 in the *Yeah, really* dialogue, nothing much hangs on whether the bank is open on Saturday, and so that the context in which the conversation finds place is a low stakes context. Nevertheless, S1 is not willing to concede that she was using the verb “know” only in an approximate fashion, and she thereby does not appear to act inappropriately.

To put the point in more general terms, we are typically aware that we are speaking only loosely when we are, or at least it does not take much to make us aware of that. However, many of our everyday knowledge ascriptions, even in low stakes contexts, simply do not appear to be made loosely. Or if they are made loosely, then for some

mysterious reason we are typically oblivious of that fact (and not easily made aware of it). To be sure, there might be some explanation of why in the case of “know” and kindred terms, we are typically oblivious of the fact that we are using them loosely (when we are). Even if there is such an explanation, that makes the current proposal much less attractive than it at first appeared to be, given that it makes an exception for “know” where the proposal appeared attractive precisely because it draws on a *general* linguistic mechanism.

Our goal in this paper is to offer yet another explanation of the purported linguistic evidence adduced by contextualists and pragmatic encroachers. Our account of the evidence is not only compatible with Bayesian decision theory, it is even—as will be seen—supported by decision-theoretic considerations. Indeed, as far as coherence with background theories goes, it is firmly rooted in important recent work in cognitive science, in particular, the development of the conceptual spaces approach briefly mentioned above. While our account has affinities with the loose talk explanation circumscribed above, it points at insights about human cognition to explain how we may be unaware that we are using “know” (etc.) in an approximate fashion when we are.

In a nutshell, the story to be told in the following is this: Criteria of identity are context-dependent and may plausibly be assumed to be co-determined by pragmatic factors; this holds for criteria of identity for properties and concepts as it does for criteria of identity for objects. So, how finely we distinguish between a concept and concepts that are almost identical to it may vary from one context to another, and may depend on pragmatic factors. In particular, how finely we distinguish between the concept of knowledge and the concept of “at least approximate knowledge” in its vicinity may depend on pragmatic factors. Thus, neither is the concept of knowledge—properly understood—pragmatically encroached, nor does “know” afford a contextualist semantics: our standards of precision for distinguishing the concept of knowledge from the nearby concept of at least approximate knowledge vary with context, which is partly determined by pragmatic factors—but this is a general phenomenon, ultimately explained in terms of contextual variation in criteria of identity. Moreover, which standards for distinguishing between concepts we employ is not something entirely under the control of our will, nor is it something we are typically aware of. So, we should not be surprised to find that, when people are using “know” and “knowledge” in an approximate way only, they are typically not aware of their doing so.

Some of the key claims in this story hinge in crucial ways on the account of identity referred to in the introduction, which, as mentioned, is embedded in the conceptual spaces program. We now first turn to a description of the basic machinery that was developed in this program and then go on to point out how knowledge can be represented within the resulting framework.

2. Conceptual spaces: knowledge as a spatial concept. The conceptual spaces program emerged out of the work of a number of psychologists and cognitive scientists. It received its first synthesis in Gärdenfors [2000], a monograph which has done much

for the systematization of the program. It is no exaggeration to say that the program is presently one of the main approaches to understanding how the human mind conceptualizes reality. The central claim of the program is that properties and concepts can be represented geometrically, as regions in conceptual spaces, where conceptual spaces are one-dimensional or multidimensional structures, equipped with a metric. Objects are mapped onto points in these spaces, and the dimensions of a space correspond to qualities relative to which objects may be compared with each other. Comparisons are made in terms of the metrics defined on the spaces. The idea is that, the closer two objects (or rather, their representations) are in a given conceptual space—closeness being measured by the metric defined on the space—the more similar these objects are in the respect corresponding to the space.

To make this less abstract, consider some actual examples of conceptual spaces. One of the simplest examples of a conceptual space is a three-dimensional space with a Euclidean metric defined on it. This space can serve to represent proximity relations between objects in the world: the closer the representations of objects are in the space, the closer the objects are in reality. Another example is auditory space, which is generally taken to be a space with two dimensions, one for pitch, and one for loudness. The closer two “objects” (in this case, sounds) are represented in the space, the more similar they sound. A third example is color space, which unarguably is the hitherto best studied conceptual space. Color space is a three-dimensional Euclidean space, with one dimension representing hue—think of the color circle—one dimension representing brightness—which ranges from white to black, through all shades of gray—and one dimension representing saturation—the intensity of the color. More complicated examples of conceptual spaces that have been described in the literature, include multidimensional shape spaces (e.g., Gärdenfors [2000, Ch. 4]) and action spaces (e.g., Gärdenfors and Warglien [2011]). Gärdenfors and Zenker [2011], [2012] present the first steps towards developing conceptual spaces specifically for the representation of physical concepts.

In the conceptual spaces approach, concepts are identified with regions of conceptual spaces. For example, the concept of redness is a region of color space, and the concept of sweetness is a region of taste space. In principle, any set of points in a space counts as a region of that space. However, empirical investigations so far warrant the conclusion that concepts corresponding to the predicates of natural language all have the feature of being representable by a *convex* set of points, that is, a set of points closed under the betweenness relation (if two points belong to the set, then so do all points that lie in between them). Gärdenfors makes a technical distinction between properties and concepts: properties are a special type of concepts, to wit, concepts that correspond to regions in spaces whose dimensions are *integral*, meaning that they represent qualities that are always experienced together; a set of integral dimensions constitutes a *domain*. Dimensions that are not integral are said to be *separable*. For example, the concept of redness is a property: we cannot experience the hue of a shade without experiencing its brightness and saturation. By contrast, the concept of being an apple is not a property: we can experience an apple’s taste without experiencing its

color or its shape. The concept of being an apple is thus represented in a space consisting of several domains (e.g., it contains both color space and taste space as separate domains).

In this paper, we aim to apply the conceptual spaces approach to the concept of knowledge. Thus we must ask: What kind of space do we need in order to represent knowledge? In particular, what are the relevant dimensions and what mathematical properties do they have (are they discrete, continuous, bounded, etc.)?

There will be some disagreement as to the number of dimensions needed for representing knowledge. In the following, we will mostly assume the old justified true belief (JTB) conception of knowledge, which naturally suggests that the requisite space is three-dimensional, with a separate dimension for each of the three component parts of knowledge. While the JTB account seems recently to have regained some of its previous popularity,⁴ most epistemologists hold that some further dimension or dimensions are to be added to prevent Gettier cases from qualifying as knowledge. Keeping the number of dimensions to three will make it easier to visualize our claims, but nothing of essence hangs on this simplification (if it is a simplification).

Plausibly, each of our three dimensions is continuous and is also bounded in at least one direction. For instance, if justification is understood in terms of the reliability of our sources and belief-forming mechanisms, it can be represented as a continuous scale ranging from 0 % (indicating complete unreliability) to 100 % (indicating complete reliability). Basically the same holds, given a coherentist understanding of justification: all popular measures of coherence make coherence a continuous variable, and most of them are bounded in at least one direction; some, like Olsson's [2002] measure and Fitelson's [2003] measure, even have a lower and an upper bound. And we can tell a similar story for modern versions of foundationalism. Acquiring a basic belief while one's senses and cognitive apparatus are functioning properly will, for most modern foundationalists, yield a maximally justified belief. More skeptically inclined foundationalists may disagree, but they will still hold that one can be maximally justified in believing simple logical truths. Furthermore, according to modern foundationalism, the support that non-basic beliefs derive from basic ones need not come from deductive inferential relations but may also derive from inductive or abductive inferential relations (or from a combination of deductive, inductive, and abductive relations). Given that both inductive and abductive inferential connections permit of continuum-many degrees,⁵ foundationalists should have no difficulty making sense of the idea that there exist continuum-many degrees of justification short of maximal justification.

Next consider truth. Starting with Popper, philosophers of science have worked on accounts of truth-closeness or verisimilitude. They have proposed both qualitative accounts—specifying what it means to say that a theory is close to the truth, or to


⁴See, e.g., Hetherington [2012], Starmans and Friedman [2012], and Turri [2012]. See Nagel, San Juan, and Mar [2013] and Starmans and Friedman [2013] for worthwhile discussion.

⁵At least this is so on all of the quantitative measures of inductive and abductive strength that one finds in the literature.

say that one theory is closer to the truth than another—and quantitative accounts, which provide a measure of how close a theory is to the truth. Given that the truth is maximally close to itself, such measures are bounded in at least one direction.⁶

And finally, belief is normally conceived of as being measurable on a scale from 0 (indicating full disbelief) to 1 (indicating full belief), with continuum-many intermediate degrees of strength of belief.

It is to be noted that these dimensions are separable. For instance, we can assess the strength of a belief without assessing its truth-closeness or its justificatory status. Thus, assuming the JTB analysis, we can represent the epistemic status of doxastic states by means of the following conceptual space consisting of three domains:



In this space, doxastic states are represented as triples of coordinates, indicating how strong the belief is, how good the grounds are for holding the belief (where we leave open the question of how exactly to understand justification), and how close to the truth the content of the belief is. Let f be the function that maps any doxastic state onto its corresponding triple of coordinates.

If we want to measure distances in epistemic status between doxastic states, we need to equip our space with some measure δ , which we will at least want to satisfy the following adequacy constraint:

(AC) $\delta(f(s), f(s^*)) < \delta(f(s), f(s'))$ iff doxastic state s^* is epistemically more similar to s than s' is similar to s .

It is worth stressing that similarity judgments concerning epistemic status need not coincide with overall similarity judgments. We may deem s more similar to s^* than to s' , say, because the propositional content of s concerns baseball and that of s^* concerns ice hockey, whereas the propositional content of s' concerns politics. Abstracting from their content, however, and just focussing on their epistemic features, s may well be more similar to s' than to s^* , perhaps because the contents of s and s' are true while that of s^* is false, or because in s and s' the contents are believed to the maximum degree whereas in s^* this is not the case.

It is not a priori how exactly δ should be defined. In principle, this metric could be determined experimentally, by asking people to judge how similar various doxastic states are to each other and deriving the metric from those judgments. (For the details of the relevant statistical technique needed for this—called “multidimensional scaling”—see Clark [1993, Appendix] and Gärdenfors [2000:21–30]; see also Decock and

⁶See Oddie [2007] for an overview of the literature and for some examples of measures of truth-closeness. Authors concerned with vagueness have also proposed accounts in which truth admits of degrees. See, for instance, Edgington [1997] and Smith [2008].




Figure 1: Three-dimensional doxastic space

Douven [2013].) For present purposes, however, the exact definition of δ is immaterial. For all we know, δ could be a so-called city block metric or it could be a Euclidean metric, to mention but the two types of metric most frequently encountered in the literature on conceptual spaces. Let $f(s) = \langle s_1, s_2, s_3 \rangle$, and similarly for $f(s^*)$ (the order of the dimensions is arbitrary). Then this is a city block metric definition of δ :

$$\delta_C(f(s), f(s^*)) := |s_1 - s_1^*| + |s_2 - s_2^*| + |s_3 - s_3^*|.$$

And as a Euclidean metric, δ would look thus:

$$\delta_E(f(s), f(s^*)) := \sqrt{(s_1 - s_1^*)^2 + (s_2 - s_2^*)^2 + (s_3 - s_3^*)^2}.$$

It might be that some dimension or dimensions matter more for our judgments of similarity *qua* epistemic status than others. If so, then this can be accounted for by adding weights to the metric; see Gärdenfors [2000:20]. But, to repeat, we remain noncommittal on the precise definition of δ and only assume that it satisfies the adequacy constraint (AC).

Although the dimensions of our space are separable, we can, for purposes of illustration, still combine them into one picture, which will enable us to visualize knowledge as a region in (what we may call) doxastic space; see Figure 1. It is not to be overlooked that this actually stands for a family of models, each model having its own specification—in terms of reliability, coherence, ...—of the justification dimension.

Where, in this model, is the concept of knowledge to be found? Sosa [2000:6] notes that “[o]ur concept of knowledge involves various dimensions each admitting a threshold,” explicitly mentioning belief (“how sure must one be?”), justification (“how much rational support is required for one’s belief?”), as well as reliability and safety; he could also have mentioned truth: how close to the truth must one’s belief be? Where one thinks the concept of knowledge is located in the above model depends precisely on which points along the individual dimensions one thinks mark the thresholds indicating sufficiency for knowledge.

Those who—like Moore [1962], Unger [1975], and Williamson [2000]—hold that knowledge requires subjective certainty of truth will think the threshold for the belief

dimension coincides with the upper endpoint of the belief scale; others—like Radford [1966], DeRose [1996], and Fantl and McGrath [2009]—will want to place this threshold a bit lower.

None of the popular theories of justification implies that the threshold for justification is at the upper bound of the justification scale. For example, reliabilists do not think that 100 % reliability is required for knowledge, nor do coherentists require perfect coherence for knowledge.

Does knowledge require truth, or will something very close to the truth do?⁷ If you believe that it takes 2 hours 4 minutes from Brussels to London by train, can you know that it takes 2 hours 4 minutes from Brussels to London even if 99 % of the services take between 2 hours 3 minutes 30 seconds and 2 hours 3 minutes and 50 seconds? Or, if you believe that the Eiffel Tower is three hundred twenty-four meters high, can you know that it is three hundred twenty-four meters high if it is 323.992 meters high?⁸ Some might argue that one should only ever believe that, for instance, the Eiffel Tower is *approximately* so-and-so high, or that it takes *approximately* 2 hours 4 minutes from Brussels to London by train. But imagine a child having read in the science section of a newspaper that new highly precise measurements have shown the Eiffel Tower to be 323.992 meters high. As the article did not go into the inherent limitations of even the highly precise method that was being used, and as the child is not aware that there are such limitations to any measurement method, she believes that the Eiffel Tower was found to be 323.992 meters high, and it would seem justifiedly so. In our view, the further question of whether she knows the Eiffel Tower to be 323.992 meters high has no obvious answer.

Be this as it may, depending on one's positions vis-à-vis the questions of whether knowledge requires certainty, and so on, one will think that, in our three-dimensional space, knowledge is best represented as something like a little cube or cuboid, or as a small rectangle, or as a line segment.⁹ Importantly, only extreme skeptics will hold that the concept of knowledge is represented by the single point where the axes of the coordinate system intersect.

We expect neither contextualists nor pragmatic encroachers to object to our spatial modelling of the concept of knowledge. However, contextualists might want to account for the apparent context-sensitivity of people's willingness to attribute knowledge in the spatial framework by claiming that the thresholds on the various dimen-

⁷As an anonymous referee rightly reminded us, knowledge is generally taken to be factive. Whether this makes knowledge incompatible with a non-maximal degree of truth is not so clear, however. For instance, Edgington [1997] can be read as suggesting that a degree of truth close to the maximum may be enough for qualifying as true in a categorical sense. Also, Cobreros et al.'s [2012] distinction between strict and tolerant truth, where the latter is defined in terms of similarity to classical truth, may help to account for our intuition that knowledge is factive even if knowledge is compatible with a non-maximal degree of truth: knowledge may be factive in the sense that it requires tolerant truth, but no strict truth.

⁸If you believe that it is 323.992 meters high, can you know that it is 323.992 meters high if it is 323.9924 meters high? And so on.

⁹If geometrically representing knowledge requires more than three dimensions, then the question is to be asked in terms of hypercubes, hyperplanes, etc.




Figure 2: Possible locations of the concept of knowledge in doxastic space

sions can vary with context (see Grimm [2011]), and that hence the location of the concept of knowledge in our space can vary with context. But this still fails to address the concern that, by having the *extension* of “knowledge” and kindred terms shift from one context to another, contextualism makes these terms come out as being linguistically odd, at best. Pragmatic encroachers, on the other hand, might insist on adding at least one other dimension to our doxastic model, to wit, a dimension representing contextual stakes. As we shall see, however, there is no need for going this way. In the following, we aim to show that the variability data that motivate this paper can be accounted for by reference to the above doxastic space—or perhaps a doxastic space containing some further non-pragmatic dimensions (for safety, or sensitivity, say)—and a general view of how we individuate concepts. The argument for this claim builds, along with the spatial approach to knowledge, on a contextual theory of identity that was presented in Douven and Decock [2010] and that we will now first briefly summarize.

3. Contextual identity. Douven and Decock’s [2010] theory of identity is primarily aimed at solving the so-called paradoxes of identity, paradoxes of coincidence—such as the famous problem case of the statue and the lump of bronze of which it is made—and paradoxes of change over time, such as the Ship of Theseus paradox and paradoxes of personal identity. On Douven and Decock’s analysis, these paradoxes show that people’s identity judgments are context-dependent and can be vague and intransitive, and thereby put pressure on the traditional view of identity as a simple and unproblematic relation. As Douven and Decock observe, there is very little unity in the solutions to these paradoxes that have been proposed so far. For example, while most popular solutions to the paradoxes of personal identity are in terms of something like psychological continuity (e.g., Parfit [1984]), one popular type of solution to the Ship of Theseus paradox is in mereological terms (e.g., Krecz [1986]), and a popular solution to the statue/lump of bronze puzzle is stated in terms of constitution (e.g., Wiggins [1980]). It is Douven and Decock’s explicit goal to offer a unified treatment of all paradoxes of identity by rethinking the concept of identity involved in those paradoxes.

In their paper, Douven and Decock argue that the concept of identity that figures in the identity claims at issue in the paradoxes of identity, and also in many other

ordinary identity claims, is not the traditional notion of identity that figures in logic textbooks, but one to be analyzed in terms of relevant similarity. Specifically, they argue that, often, when people claim that two things are identical, what these people mean is that the things are highly similar in all relevant respects. Douven and Decock cite psychological research showing that, in comparing two or more objects, people rarely take into account all the respects in which the objects might be found to be similar to each other but rather attend to a proper subset of those respects, and that which subset people attend to may vary from one context to another. Thus, if identity is often interpreted as high similarity in all relevant respects, and if it depends on context which respects are relevant, it is a matter of course that identity judgments can vary with context. For instance, a statue and the lump of bronze of which it is made may be judged to be identical in one context and non-identical in another. Moreover, while two objects either are or are not identical to each other given the logical notion of identity—which excludes the possibility of vague identity—whether two objects are highly similar to one another can be vague; so in the sense envisaged by Douven and Decock, things can be vaguely identical. And finally, since similarity is not generally transitive—it can be that a is highly similar to b , b highly similar to c , yet a not highly similar to c —it is, on the proposed concept of identity, also clear how people’s identity judgments can fail to be transitive.

Douven and Decock use the conceptual spaces framework to make precise their proposed alternative concept of identity, notably, to make precise the notion of being highly similar in all contextually relevant respects. In a first step, they identify a contextually relevant respect with a conceptual space that is being activated in the given context. Concomitantly, they associate with each context C the set S_C of conceptual spaces that contains a space for each respect that is relevant in C . For instance, if shape is a relevant respect of comparison in C , then S_C contains a shape space; if color is relevant in C , S_C contains color space; and so on. Then, where $d_r(\cdot, \cdot)$ is the distance function associated with space $r \in S_C$, and where $t_C^r \geq 0$ is a threshold value associated with the same space, they formally define the alternative concept of identity as follows:

$$(ID) \quad \text{Id}_C(a, b) \iff \forall r \in S_C : d_r(a, b) \leq t_C^r.$$

The thresholds associated with the spaces are meant to capture the requirement that things that are identical in the sense at issue be *highly* similar, or similar *enough*, in all relevant respects. The superscripted r and subscripted C indicate that, on this proposal, what counts as highly similar may vary both with respect and with context.

Douven and Decock note that, in their account, context shifts can be brought about by a variety of changes. A respect previously deemed irrelevant may become relevant, or the threshold or thresholds associated with one or more of the activated spaces may change, or, if the metric is weighted (see page 10), the weights may change, or—and this may hold for the majority of context shifts—a number of such changes may occur simultaneously. As Douven and Decock also note, however, determining the finer details of such changes is a matter of empirical investigation.

Furthermore, although, as we said, the above account of identity was in the first instance proposed with an eye toward solving the paradoxes of identity, which all involve identity statements relating material objects, Douven and Decock [2010:72 f] point out that the account is not limited to such statements and applies equally well to identity statements relating properties or indeed anything that can be represented as a point or set of points in a conceptual space. For just as one can measure distances between points in a conceptual space, one can measure distances between sets of points in a conceptual space. The best known metric for this purpose is the so-called Hausdorff metric, according to which the distance between sets Φ and Ψ equals

$$\max\{\sup_{x \in \Phi} \inf_{y \in \Psi} \delta_S(x, y), \sup_{y \in \Psi} \inf_{x \in \Phi} \delta_S(x, y)\}.$$

We might thus, in some contexts, affirm that ochre and golden-yellow are identical colors, while denying this in other contexts, in which we want to distinguish more finely between colors: the Hausdorff distance between the region representing ochre in color space might be below the relevant threshold value in one context but not in the other. This fact will be exploited in the next section to explain why our knowledge attributions may vary with context.

4. An alternative explanation of the variability data. Our aim is to explain the apparent context-dependence of people’s willingness to attribute knowledge to others or themselves and to do so in a way that is compatible with “knowledge” having a fixed extension across contexts, and a meaning that does not involve reference to pragmatic concerns. Our basic idea relies on the machinery introduced in the previous sections and is simple enough: often, we do not distinguish between the concept of knowing and that of at least approximately knowing, and whether we do distinguish between the two may well depend on what is at stake in a given context.

To explain this in detail, we begin by noting that, in our geometric framework, the concept of *at least approximate knowledge* can be defined formally as the set of points

$$\{p \mid \exists k \in \mathcal{K} : \delta(p, k) \approx 0\},$$

where \mathcal{K} is the set of points that lie in whichever region of our three-dimensional space that represents the concept of knowledge. Accordingly, if we let s_φ^P denote person P ’s doxastic attitude toward φ , then we can say that P is *at least close to knowing* φ , or that P *knows* φ *at least approximately* iff

$$f(s_\varphi^P) \in \{p \mid \exists k \in \mathcal{K} : \delta(p, k) \approx 0\}.$$

The \approx relation can be spelled out in terms of a threshold θ , in that $a \approx b$ can be defined to hold iff $|a - b| \leq \theta$. By way of general fact about \approx , or about the expression “approximately equal,” the value of θ may be vague and may itself depend on context. To keep things simple, we abstract from both the vagueness and the context-dependence involved in pinpointing the value of θ and rather focus on the context-dependence of

people's distinguishing between the concept of knowledge and that of at least approximate knowledge.

Just as, in many contexts, we may deem a statue and the lump of bronze of which it is made to be identical, we may, in many contexts, identify at least approximate knowledge with knowledge: the two concepts are too similar in all contextually relevant respects to be worth distinguishing between—too similar, that is, given the standards operative in the given context. And just as a shift in context may have us distinguish between the statue and the lump of bronze—for instance, because we are made to attend to some modal difference between the two—a shift in context may make us distinguish between knowledge and at least approximate knowledge, for instance, because we are made aware of what is at stake in believing correctly for the putative knower. This means that a doxastic state $a \notin \mathcal{K}$ whose representation in doxastic space is very close to some doxastic state $b \in \mathcal{K}$ may in some contexts qualify as a state of knowledge while in others it does not.

This is no more than an application of Douven and Decock's account of identity to the concepts of knowledge and at least approximate knowledge. But one may wonder how, apart from terminology, this is different from the “varying standards of precision” response that we rejected in Section 1. Independently, one may wonder where the putative connection with stakes is supposed to come from. As intimated, in their paper on identity, Douven and Decock relegate a detailed investigation of the question of what occasions context-shifts to empirical investigation. In any case, they do not explicitly relate their account to contextually varying stakes.

As to the first question, there are undeniable parallels between the loose talking proposal and ours. Both proposals agree that neither are there shifts across contexts in what “know” refers to nor does someone's knowledge vary according to what is at stake for her. Specifically for our proposal, “knowledge” designates the region of doxastic space fixed by whatever the precise values of the thresholds on the justification, belief, and truth dimensions are, where these thresholds are *not* context-dependent. Any contextual variation derives from the fact that we do not always care to distinguish that region from a slightly more inclusive region designating the concept of at least approximate knowledge. This could be said to be a form of loose thinking. And for the same reason that we do not always aim to speak strictly—sometimes it just is not worth the extra effort, as Krifka and others have noted—we do not always aim at the greatest precision in our thinking.

In spite of these parallels, however, our proposal has an important advantage as compared to the loose talking proposal. Recall the problem facing the latter: whereas it is typically clear to us that we are speaking loosely when we are, or at least not much is needed to make this clear, we do *not* typically feel that we are using “know” or “knowledge” loosely, whatever the stakes at issue. This asymmetry, we claim, does not arise for our present proposal. Quite the contrary, as comparison of the following dialogues suggests:

Statue

- S1: The statue and the lump of bronze are one and the same object.
S2: Really?
S1: Yeah, really!
S2: But what if I'd use some heavy machinery to flatten the lump of bronze? Thereby the statue would cease to exist, but not the lump of bronze.
S1: I guess you are right. The two are not identical.

Bank

- S1: We could go to the bank tomorrow instead of today.
S2: Tomorrow is Saturday—do you know whether the bank is open then?
S1: Yes, it is.
S2: Do you really know that?
S1: Sure!
S2: But banks sometimes change their opening hours for Saturdays.
S1: I guess you are right. I don't know that the bank will be open tomorrow.

In *Statue*, it is only after a modal dimension along which the objects may be compared has been brought to S1's attention that S1 is willing to reconsider. Equally in *Bank*, it is only after an error possibility has been brought to S1's attention that S1 is willing to reconsider. In neither case is a simple prompting—"Really?"—enough to give S1 pause, nor, we submit, should it do so.

This phenomenon relates to a more general point made in Douven and Decock [2010:76]. As they note there, in every particular context, a judgment to the effect that two things are identical in the contextual sense may plausibly be assumed to "feel" like a judgment as to these things being identical in the traditional philosophical sense. One reason for this is that—Douven and Decock argue—people may be expected to read the restricted quantifier "for all contextually relevant respects," which occurs in (ID), as an unrestricted one, given that all but the respects quantified over are being ignored in the given context. Moreover, in all contexts identity requires things to be highly similar in the relevant respects. And while the standards for qualifying as highly similar may vary across contexts, this variation is unrelated to the loose versus strict talking contrast: when we call a road flat, we do not consider ourselves to be talking loosely even if we are willing to concede that the road does not qualify as flat by the standards for flatness assumed in a physics lab, and we stop calling the road flat, once these higher standards have been mentioned in the conversation. Applied to the above dialogues, the point is that, "from the inside," it is as though we are speaking strictly when we first identify the statue with the lump of bronze, as well as when we first claim to know that the bank is open on Saturday; it takes some effort to make us alter our judgments on these matters. All this is different when we assert that it takes two hours from Brussels to London; then, as mentioned, we normally realize fully well that we are speaking loosely.

We now turn to what may be the crucial question that will have to be answered if we want to claim that our account is able to explain the variability data, to wit, the question of why we should want to distinguish more carefully between knowledge and

approximate knowledge in high stakes contexts than in low stakes contexts. To begin answering this question, we compare two decision-theoretic situations, a low stakes case and a high stakes case, where the agent faces the choice to pay a check on Friday or to wait till Saturday and then pay the check. The table below gives the pay-off matrix for these options in the low stakes case, where the relevant hypotheses are that the bank is open and closed, respectively, on Saturday:

	Open on Saturday	Closed on Saturday
Pay Friday	2	2
Pay Saturday	3	1

Suppose that this matrix represents a given agent's utilities and that this agent's degree of belief that the bank is open on Saturday equals .9. Then her expected utility for paying on Friday is 2 and that for paying on Saturday is 2.8. So, assuming the agent to be rational (in the sense of Bayesian decision theory), she will choose to pay on Saturday. Now suppose that her degree of belief that the bank is open on Saturday is not .9 but 1. Then her expected utility for paying on Friday is again 2, but her expected utility for paying on Saturday has gone up to 3. Nevertheless, this will make no difference to her choice: she will still choose to pay on Saturday.

The pay-off matrix for the high stakes case is given by this table:

	Open on Saturday	Closed on Saturday
Pay Friday	20	20
Pay Saturday	21	1

If, given these utilities, the agent's degree of belief that the bank is open on Saturday equals 1, then her expected utility for paying on Friday is 20 and that for paying on Saturday is 21, and so the agent will again choose to pay on Saturday. However, this time, if her degree of belief that the bank is open on Saturday equals .9, her expected utility for paying on Friday is still 20 but that for paying on Saturday is 19. Thus, in this situation, the rational agent will choose to pay on Friday.

Now assume that knowledge requires subjective certainty, or at least a degree of belief exceedingly close to 1. Insofar as approximate knowledge leaves open the question of whether someone is or is not (very close to being) entirely sure of what she is said to know, in the low stakes case it does not matter to the agent's choice nor to our explanation of her choice whether or not she knows, strictly speaking, that the bank is open on Saturday. However, in the high stakes case it does. This shows that what is at stake for a putative knower may determine whether, in attributing (or not attributing) knowledge to her, we should distinguish more or less finely between knowledge and at least approximate knowledge.

To be sure, this is only one example. Given that, in spite of all the talk about high and low stakes in the epistemological literature of the past ten years, no one has tried to state in even remotely general terms what high and low stakes amount to, it is not straightforward to demonstrate in an entirely general fashion that in high stakes contexts it is more important to distinguish between knowing and at least approximate knowing than it is in low stake contexts. But if the cases that are unexceptionally discussed in connection with contextualism and pragmatic encroachment—the already cited Bank Cases, after which the above scenario was modelled, as well as the Airport Cases—capture at least the gist of the high stakes/low stakes contrast, then we can go well beyond the above illustration in showing how high stakes may require making finer distinctions between knowledge and at least approximate knowledge.

Consider the utility matrix given in the table below. Assume that utility x is related to the main goal of the relevant agent and represents what is at stake for her in achieving this goal. Utility y represents a secondary motive for the agent, related to matters of convenience or avoidance of small inconveniences.¹⁰ We assume that achieving the main goal is always at least slightly more important than matters of convenience; hence, $x > y$. In particular, for the low stakes case $x = x_L \gtrsim y$, and for the high stakes case $x = x_H \gg y$.¹¹

	h_1	h_2
a_1	x	x
a_2	$x + y$	y

The agent can choose between two available actions, a_1 and a_2 . Option a_1 is safe in that it ensures that the agent will achieve her main goal; option a_2 does not come with this guarantee, but it does ensure achieving the secondary goal. Hypothesis h_1 describes a situation in which the agent will always achieve her main goal independently of the action she chooses; hypothesis h_2 describes a situation in which the agent does not achieve her main goal if she opts for action a_2 ; these hypotheses are taken to be jointly exhaustive. The agent does not know for certain which of the two hypotheses is true: she holds a non-extreme degree of belief of p in h_1 and a degree of belief of $1 - p$ in h_2 .¹²

¹⁰We have defined the utilities such that they are always positive. This choice is inconsequential: only the relative utilities matter, since they determine the differences between expected utilities of various actions (i.e., the absolute utilities are on an interval scale rather than a ratio scale).

¹¹In the example in the text above, we had that $x_L = 2$, $x_H = 20$, and $y = 1$.

¹²It should be easy to see how this schematic situation fits both the Bank Cases and the Airport Cases. For the Bank Cases, let h_1 be the hypothesis that the bank is open on Saturday, and let $h_2 = \neg h_1$. Furthermore, let a_1 be the option of going to the bank on Friday; a_2 the option of going to the bank on Saturday; x the utility of going to the bank while it is open; and y the utility of avoiding the long queues on Friday. For the Airport Cases, let h_1 be the hypothesis that the intended flight has a layover in Chicago, with again $h_2 = \neg h_1$. Let a_1 be the option of contacting the airline and rescheduling the flight if necessary; a_2 the option of taking the intended flight; x the utility of having a layover in Chicago; and y the utility of avoiding the hassle of contacting the airline.

To state our earlier point in more general terms, then, note that the agent's expected utility for a_1 , $EU(a_1)$, equals $px + (1 - p)x = x$, while her expected utility for a_2 , $EU(a_2)$, equals $p(x + y) + (1 - p)y = y + px$. Both for low and high stakes, we thereby know the preferred action in the case of extreme degrees of belief. For $p = 0$, $EU(a_1) = x > y = EU(a_2)$, so that in this case the rational agent will opt for a_1 , independently of the stakes.¹³ For $p = 1$, $EU(a_1) = x < x + y = EU(a_2)$, so in this case the rational agent will opt for a_2 , again independently of the stakes.¹⁴

However, for non-extreme degrees of belief, $p \in (0, 1)$, the agent's preferred action switches as a function of p ; the change in preference between a_1 and a_2 occurs when $EU(a_1) = EU(a_2)$, that is, when $p = 1 - y/x$. For low stakes, the agent prefers a_1 if $p \in [0, 1 - y/x_L)$ and opts for a_2 if $p \in (1 - y/x_L, 1]$. For high stakes, the agent prefers a_1 if $p \in [0, 1 - y/x_H)$ and opts for a_2 if $p \in (1 - y/x_H, 1]$. Hence, if the agent's degree of belief p is in the interval $(1 - y/x_L, 1 - y/x_H)$, then she will rationally opt for a_1 in the case of high stakes, while it is equally rational for her to opt for a_2 in the case of low stakes. See Figure 3 for a graphical representation of this result.

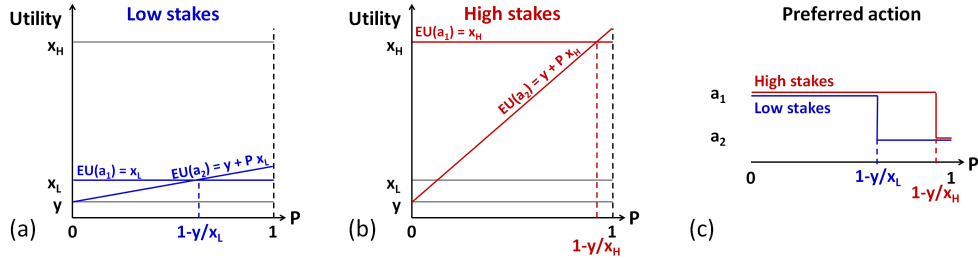


Figure 3: In these graphs, the horizontal axis represents the degree of belief of an agent in hypothesis 1 (p). On the vertical axis in panels (a) and (b), there are three relevant utility levels: the primary utility in a case with low stakes (x_L), the primary utility in a case with high stakes (x_H), and a constant secondary utility (y). (a) In a case with low stakes, $EU(a_2) > EU(a_1)$ for degrees of belief well below unity. (b) In a case with high stakes, $EU(a_2) > EU(a_1)$ only for degrees of belief that are very close to unity. (c) For degrees of belief in the interval $(1 - y/x_L, 1 - y/x_H)$, the preferred action depends on the stakes.

This analysis shows that the higher the stakes are, the higher the agent's degree of belief in h_1 has to be for her preference to switch. As a result, not distinguishing between knowledge and at least approximate knowledge is more likely to be consequential in a high stakes case than it is in a low stakes case: in the former case, it is more likely to influence the agent's decision, and the more so, the higher the stakes

¹³ Applied to the Bank Cases: if the agent is completely certain that the bank will be closed on Saturday, she will go there on Friday, which is her only option to achieve her main goal.

¹⁴ Applied to the Bank Cases again: if the agent is completely certain that the bank will be open on Saturday, she will go there on Saturday, which allows her to achieve her main goal while avoiding the queues on Friday.

are. Look again at Figure 3: we are assuming that approximate knowledge is compatible with having a degree of belief unequal to 1 or a value very close to 1, and that it only requires a degree of belief above a certain threshold. As panels (a) and (b) of the figure help us to appreciate, this threshold is more likely to be found in the interval $(1 - \gamma/x_L, 1]$ than in the interval $(1 - \gamma/x_H, 1]$, and the more likely, the smaller x_L is and the larger x_H is.

To be sure, this explanation assumes standard decision theory. But, as pointed out in Section 1, that assumption is pretty innocuous. Indeed, we argued there that one of the standard explanations of the phenomena at issue in this paper is marred by the fact that it is at odds precisely with standard decision theory. By contrast, our explanation is not only consistent with standard decision theory, it is, as we just saw, supported by that theory.¹⁵

Also note that, while it may be harder to make parallel points regarding justification and truth in an equally formal fashion, it is not hard at all to imagine that sometimes, but not always, it matters whether a person is justified or only close to being justified in believing a proposition, and equally that sometimes, though not always, it matters whether someone's belief in a proposition is true or whether it is only approximately true. If either of these things matter, that is another reason to make fine distinctions between the concept of knowledge and that of at least approximate knowledge.

By way of further support, we note that our proposal also explains why, as we said earlier, philosophers tend to disagree about whether knowledge requires certainty, as well as why it is so hard to determine whether we can know, strictly speaking, that the Eiffel Tower is three hundred twenty-four meters high if it is 323.992 meters high. In trying to answer these questions, we bring to the task intuitions that are informed by our everyday usage of "know" and related words. If, in many everyday contexts, we use these words without discriminating too closely between knowledge and at least approximate knowledge, where, as explained, phenomenologically in such contexts we seem to be speaking as strictly as we are in contexts in which we do discriminate between knowledge and approximate knowledge, we should expect the kind of confusion we find about the above questions even if knowledge does require certainty and knowing that the Eiffel Tower is three hundred twenty-four meters high is incompatible with the Eiffel Tower not being precisely 324.00 meters high.

There might be further evidence of the same sort. Above we said that, while objects may in principle be compared with each other along indefinitely many dimensions, we usually focus on only a relatively small number of those dimensions. One way in which a context shift can occur is, then, that a dimension previously unattended to becomes salient. Arguably, there are not indefinitely many dimensions along which the concept of knowledge and that of at least approximate knowledge could be mean-

¹⁵As an anonymous referee observed, variants of DeRose's Bank Cases could be used to argue for the context-sensitivity of attributions of certainty. The referee also observed that our decision-theoretic explanation of the context-sensitivity of knowledge attributions would generalize swiftly to an explanation of the context-sensitivity of certainty attributions.

ingfully compared with one another; even if there are more than the three we have been assuming in this paper, there are at most a handful. Context shifts relevant to the explanation of the variability data will mostly be due to a change in the threshold value for counting two regions in doxastic space as being “similar enough.” Still, there is at least a theoretical possibility that some of the dimensions along which we may compare knowledge and at least approximate knowledge go unattended for in some contexts. And perhaps this is even a possibility to reckon with in practice. For the supposition that, in attributing knowledge, we do not always attend to all dimensions of doxastic space would go some way toward explaining why a number of theorists have the intuition that knowledge is not factive (e.g., Hazlett [2010]) or the intuition that knowledge does not require justification (e.g., Sartwell [1992] and, to some extent, Hetherington [2001] and Goldman and Olsson [2009]). Such intuitions might have been nourished by usage of the words “know” and “knowledge” in contexts in which the truth, respectively justification, dimension are ignored.

Our main purpose in this paper has been to explain the data about knowledge attributions in terms of a pre-existing framework and a pre-existing account of identity. We hope to have shown that, jointly, these elements correctly predict the data, and that they do so without assuming that knowledge is context-sensitive or suffering from some kind of encroachment, pragmatic or otherwise. As far as the results of our paper go, the concept of knowledge may be the same region in doxastic space across all contexts.¹⁶

References

- Bach, K. [2005] “The Emperor’s New ‘Knows,’” in G. Preyer and G. Peter (eds.) *Contextualism in Philosophy*, Oxford: Oxford University Press, pp. 51–89.
- Clark, A. [1993] *Sensory Qualities*, Oxford: Clarendon Press.
- Cobrerros, P., Égré, P., Ripley, D., and van Rooij, R. [2012] “Tolerant, Classical, Strict,” *Journal of Philosophical Logic* 41:347–385.
- Cohen, S. [1987] “Knowledge, Context, and Social Standards,” *Synthese* 73:3–26.
- Cohen, S. [1988] “How to Be a Fallibilist,” *Philosophical Perspectives* 2:91–123.
- Cohen, S. [1999] “Contextualism, Skepticism, and the Structure of Reasons,” *Philosophical Perspectives* 13:57–89.
- Davis, W. [2007] “Knowledge Claims and Context: Loose Use,” *Philosophical Studies* 132:395–438.
- Decock, L. and Douven, I. [2013] “What Is Graded Membership?” *Noûs*, in press.
- DeRose, K. [1992] “Contextualism and Knowledge Attributions,” *Philosophy and Phenomenological Research* 52:913–929.
- DeRose, K. [1995] “Solving the Sceptical Puzzle,” *Philosophical Review* 104:1–52.

¹⁶We are greatly indebted to René van Woudenberg and to an anonymous referee for this journal for very helpful comments on previous versions of this paper.

- DeRose, K. [1996] "Knowledge, Assertion, and Lotteries," *Australasian Journal of Philosophy* 74:568–580.
- Douven, I. [2004] "The Context-Insensitivity of 'Knowing More' and 'Knowing Better,'" *Canadian Journal of Philosophy* 34:313–326.
- Douven, I. [2008a] "Underdetermination," in S. Psillos and M. Curd (eds.) *The Routledge Companion to the Philosophy of Science*, London: Routledge, pp. 292–301.
- Douven, I. [2008b] "Knowledge and Practical Reasoning," *Dialectica* 62:101–118.
- Douven, I. and Decock, L. [2010] "Identity and Similarity," *Philosophical Studies* 151:59–78.
- Edgington, D. [1997] "Vagueness by Degrees," in R. Keefe and P. Smith (eds.) *Vagueness: A Reader*, Cambridge MA: MIT Press, pp. 294–316.
- Fantl, J. and McGrath, M. [2009] *Knowledge in an Uncertain World*, Oxford: Oxford University Press.
- Feldman, R. [2001] "Skeptical Problems, Contextualist Solutions," *Philosophical Studies* 103:61–85.
- Fitelson, B. [2003] "A Probabilistic Theory of Coherence," *Analysis* 63:194–199.
- Gärdenfors, P. [2000] *Conceptual Spaces*, Cambridge MA: Bradford.
- Gärdenfors, P. [2007] "Representing Actions and Functional Properties in Conceptual Spaces," in T. Ziemke, J. Zlatev, and R. M. Frank (eds.) *Body, Language and Mind* (Vol. 1), Berlin: De Gruyter, pp. 167–195.
- Gärdenfors, P. and Warglien, M. [2011] "Using Concept Spaces to Model Actions and Events," *Journal of Semantics*, in press.
- Gärdenfors, P. and Zenker, F. [2011] "Using Conceptual Spaces to Model the Dynamics of Empirical Theories," in E. J. Olsson and S. Enqvist (eds.) *Belief Revision Meets Philosophy of Science*, Berlin: Springer, pp. 137–153.
- Gärdenfors, P. and Zenker, F. [2012] "Theory Change as Dimensional Change: Conceptual Spaces Applied to the Dynamics of Empirical Theories," *Synthese*, in press.
- Goldman, A. and Olsson, E. J. [2009] "Reliabilism and the Value of Knowledge," in A. Haddock, A. Millar, and D. Pritchard (eds.) *Epistemic Value*, Oxford: Oxford University Press, pp. 19–41.
- Grimm, S. R. [2011] "On Intellectualism in Epistemology," *Mind* 120:705–733.
- Hawthorne, J. [2004] *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Hazlett, A. [2010] "The Myth of Factive Verbs," *Philosophy and Phenomenological Research* 80:497–522.
- Hetherington, S. [2001] *Good Knowledge, Bad Knowledge*, Oxford: Oxford University Press.
- Hetherington, S. [2012] "There Can Be Lucky Knowledge," in M. Steup, J. Turri, and E. Sosa (eds.) *Contemporary Debates in Epistemology*, Oxford: Wiley–Blackwell, in press.
- Kennedy, C. [2007] "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives," *Linguistics and Philosophy* 30:1–45.

- Klein, P. [2000] "Contextualism and the Real Nature of Academic Skepticism," *Philosophical Issues* 10:108–116.
- Kornblith, H. [2000] "The Contextualist Evasion of Epistemology," *Philosophical Issues* 10:24–32.
- Krecz, C. [1986] "Parts and Pieces," *Philosophy and Phenomenological Research* 46:381–400.
- Krifka, M. [2002] "Be Brief and Be Vague," in D. Restle and D. Zaefferer (eds.) *Sounds and Systems: Studies in Structure and Change*, Berlin: Mouton de Gruyter, pp. 439–458.
- Lewis, D. [1979] "Scorekeeping in a Language Game," *Journal of Philosophical Logic* 8:339–359.
- Lewis, D. [1981] "Causal Decision Theory," *Australasian Journal of Philosophy* 59:5–30.
- Lewis, D. [1996] "Elusive Knowledge," *Australasian Journal of Philosophy* 74:549–567.
- Moore, G. E. [1962] "Certainty," in his *Philosophical Papers*, New York: Collier Books, pp. 226–251.
- Nagel, J., San Juan, V., and Mar, R. [2013] "Lay Denial of Knowledge for Justified True Beliefs," *Cognition*, in press.
- Oakley, I. T. [2001] "A Skeptic's Reply to Lewisian Contextualism," *Canadian Journal of Philosophy* 31:309–332.
- Oddie, G. [2007] "Truthlikeness," *The Stanford Encyclopedia of Philosophy* (Fall 2008 Edition), available at <http://plato.stanford.edu/entries/truthlikeness/>.
- Olsson, E. J. [2002] "What is the Problem of Coherence and Truth?" *Journal of Philosophy* 94:246–272.
- Parfit, D. [1984] *Reasons and Persons*, Oxford: Oxford University Press.
- Radford, C. [1966] "Knowledge—By Examples," *Analysis* 27:1–11.
- Rysiew, P. [2001] "The Context-Sensitivity of Knowledge Attributions," *Noûs* 35:477–514.
- Sartwell, C. [1992] "Why Knowledge Is Merely True Belief," *Journal of Philosophy* 89:167–180.
- Schiffer, S. [1996] "Contextualist Solutions to Skepticism," *Proceedings of the Aristotelian Society* 96:317–333.
- Smith, N. J. J. [2008] *Vagueness and Degrees of Truth*, Oxford: Oxford University Press.
- Sosa, E. [2000] "Skepticism and Contextualism," *Philosophical Issues* 10:1–18.
- Stanley, J. [2005] *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Starmans, C. and Friedman, O. [2012] "The Folk Conception of Knowledge," *Cognition* 124:272–283.
- Starmans, C. and Friedman, O. [2013] "Taking 'Know' for an Answer: A Reply to Nagel, San Juan, and Mar," *Cognition*, in press.
- Turri, J. [2012] "Is Knowledge Justified True Belief?" *Synthese* 184:247–259.

- Unger, P. [1975] *Ignorance: A Case for Scepticism*, Oxford: Clarendon Press.
- van Deemter, K. [2010] *Not Exactly*, Oxford: Oxford University Press.
- Weirich, P. [2004] *Realistic Decision Theory: Rules for Nonideal Agents in Nonideal Circumstances*, Oxford: Oxford University Press.
- Wiggins, D. [1980] *Sameness and Substance*, Oxford: Blackwell.
- Williamson, T. [2000] *Knowledge and Its Limits*, Oxford: Oxford University Press.