

## Possibilities regained: neo-Lewisian contextualism and ordinary life

Mario Piazza<sup>1</sup> · Nevia Dolcini<sup>2</sup>

Received: 16 December 2014 / Accepted: 29 October 2015 / Published online: 17 November 2015  
© Springer Science+Business Media Dordrecht 2015

**Abstract** According to David Lewis, the predicate ‘knows’ is context-sensitive in the sense that its truth conditions vary across conversational contexts, which stretch or compress the domain of error possibilities to be eliminated by the subject’s evidence (Lewis, *Aust J Philos* 74:549–567, 1996; Lewis, *J Philos Log* 8:339–359, 1979). Our concern in this paper is to thematize, assess, and overcome within a neo-Lewisian contextualist project two important mismatches between our use of ‘know’ in ordinary life and the use of ‘know’ by ‘Lewisian’ ordinary speakers. The first mismatch is that Lewisian contextualism still *overgenerates* the error possibilities which cannot be ignored in a given context, since it is oblivious to the distinction between ‘invented’ and ‘discovered’ possibilities. The second mismatch is a full-scale one: an adequate account of knowledge attribution is not exhausted by the subject’s *negative* capacity of pruning branches off the tree of counterpossibilities. We therefore introduce a new vector of value, which explains how ‘know’ comes in degrees: the satisfaction of ‘know better’ is made to depend on the capacity of *imagining* (actualized) possibilities connected in a relevant way with the subject’s (true) beliefs.

**Keywords** Epistemology · Knowledge attribution · Epistemic contextualism · David Lewis · Epistemic modals · Imagination

---

✉ Mario Piazza  
m.piazza@unich.it  
Nevia Dolcini  
ndolcini@umac.mo

<sup>1</sup> Department of Philosophy, University of Chieti-Pescara, Chieti, Italy

<sup>2</sup> Philosophy and Religious Studies Programme, Faculty of Arts and Humanities, University of Macau, Macau, China

## 1 Elusive abundance

What kind of evaluation do we perform when we inquire as to whether someone knows that something is so? Lewis's answer in 'Elusive Knowledge' is that this evaluation, or knowledge attribution, is a phenomenon both *eminently modal* and *eminently contextual*, that is shaped in subtle ways by human interactions in language. The predicate 'knows' is context-sensitive in the sense that its truth conditions vary across conversational contexts in which the predicate is employed, since these conversational contexts circumscribe the domain of error possibilities to be eliminated by the putative knower's evidence (Lewis 1996, 1979). Thus, as held by any brand of epistemic contextualism, a sentence of the form 'S knows that *p*' (even when *p* is expressed by an indexical-free statement) is not truth-evaluable as it stands: it may express a true proposition as uttered in one context, and a false one as uttered in another.<sup>1</sup> One common, and at least initially attractive, way to convey the core idea of epistemic contextualism is to say that knowledge ascriptions are context-sensitive in roughly the same way as tallness or elegance ascriptions are: a tall jockey may be a lilliputian basketball player, and an chic dress in a philosophy seminar may be adequately described as ordinary in the Academy Awards.<sup>2</sup>

The germ of the contextualist insight can be found in a famous passage from René Descartes's *Reply to Objections II* (1641). Take an atheist geometer and the proposition *p* [= the sum of the three angles of a triangle is equal to 180°]. In a mathematical context, it is quite natural and correct to ascribe knowledge (*cognitio*) that *p* to the man, by supposing that he has *proved* that *p*. Yet, it is incorrect to claim that he has knowledge (*scientia*) that *p* in a more demanding philosophical context, whereby one needs to acknowledge the existence of God in order to rule out the possibility of being deceived by an evil demon, *inter alia*, about *p* itself (Descartes 1984, p. 101).<sup>3</sup>

The general concern of 'Elusive Knowledge' is the problem of eschewing the Scylla of fallibilism and the Charybdis of skepticism, which flourishes under an across-the-board infallibilism.

*Scylla:*

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 possibility of error, just *sounds* contradictory (Lewis 1996, p. 549).

<sup>1</sup> There are several proposals falling under the heading of epistemic contextualism. We take it as primarily a thesis about knowledge attribution, which retains a measure of autonomy from any specific account of knowledge. Other leading contextualists understand the variation of knowledge attributions in terms of a scale of (contextually provided) epistemic strength (DeRose 1992; Cohen 1987). Throughout the paper, by 'contextualism' ('contextualist') we mean 'Lewisian contextualism (contextualist)'.

<sup>2</sup> Lewis's favourite example of gradable adjective is 'flat' (as opposed to 'bumpy'). The example is borrowed from (Unger 1975), who nevertheless takes the term as 'absolute' in the sense that it is inconsistent to say about two flat objects that one is flatter than the other.

<sup>3</sup> Of course, this does not make Descartes a contextualist; and indeed, as an invariantist he employs two distinct terms for knowledge—*cognitio* and *scientia*—in the mentioned passage.

However, the above passage gives no explanation of why speaking of fallible knowledge has the ring of an oxymoron. This seems to be due to the intuition that to attribute knowledge to a subject is to say something *positive* about her epistemic capacities; however, such an intuition conflicts with the case of a subject who is granted to be unable to rule out a certain possibility of error. This means that, under fallibilism, the speaker's utterances of knowledge attributions cannot be understood as expressing genuine *approval*.<sup>4</sup>

*Charybdis*: typically, the radical skeptic exploits the infallibilist principle that  $S$ 's knowledge that  $p$  requires  $S$ 's elimination of *every* possibility of error concerning  $p$ . As corollary, we can know nothing, or next to nothing, about the external world, so that virtually all our attributions of empirical knowledge are false.

According to Lewis, to “dodge the choice” between fallibilism and skepticism, one has to constrain contextually the infallibilist principle so that ‘every possibility’ is restricted to cover just the domain of possibilities which are relevant to knowledge attributions in a particular conversational context. In fact, quantifier restrictions over possibilities “come and go with the pragmatic wind”, as Lewis writes elsewhere (Lewis 1986, p. 164); that is, ‘every possibility’ responds to the pragmatics of quantifier restrictions in natural language.

A metalinguistic (and disjunctive) version of Lewis's definition of knowledge is as follows:

(Def<sub>c</sub>)  $S$  satisfies ‘knows that  $p$ ’ in context  $c$  iff every  $\neg p$ -possibility is either eliminated by  $S$ 's evidence, or it is properly ignored in  $c$ .<sup>5</sup>

Then he articulates a set of rules under which the domain of the properly ignored (or relevant) possibilities may be expanded and contracted from one conversational context to another. These rules both determine the truth-value of knowledge attributions and structure the interaction between evaluator, or knowledge attributor, and evaluated.

Lewis's contextualist solution delivers bad news, good news, and a general worry. The bad news, needless to say, is that ‘knows’ is *elusive*: it evaporates as soon as we project ourselves in the context of epistemological reflection within which skeptical hypotheses become relevant and “our paranoid fantasies rip” (Lewis 1996 p. 550). Massive use of possibilities is endemic to all sub-disciplines of philosophy. But, in particular, philosophers have always been struck by the ease with which the audience in the epistemology classroom can be brought to appreciate the overextension of the space of possibilities: evil demons, brains in a vat, Matrices, fully comprehensive dreams, five-minute-old world, and the like. If the contents of the skeptic's fantastical stories were actualized, the world would still appear to us just as it does. And such an evidential parity is the obvious reason why skeptical possibilities stay uneliminated (and ineliminable), and knowledge is the sort of thing that is ‘destroyed’ in a philosophical context, which, according to Lewis, is inevitably a skeptical one.

<sup>4</sup> For a discussion of the so-called concessive knowledge attribution in fallibilism, see Dodd (2010).

<sup>5</sup> Lewis's original definition is not metalinguistic and uses an except-clause which he calls ‘*sotto voce* proviso’.  $S$ 's evidence at a given time contains the totality of  $S$ 's current perceptual experiences and memory states; a possibility  $w$  is ‘eliminated’ by  $S$ 's evidence iff  $S$ 's entire perceptual experience and memory contains information incompatible with the  $w$ 's obtaining.

The good news is that outside the philosophy classroom we know a great deal, and this relatively easy and diffuse satisfaction of ‘knows’ accommodates pre-philosophical data as well as intuitions of a Moorean bent. To wit, in ordinary contexts ‘knows’ is *abundant*, rendering the title of Lewis’s paper somewhat misleading, or at least partial with respect to the intended contextualist moral. Speakers typically are right, or innocent until proven guilty, when they assert ‘I know that *p*’ in the various everyday conversational settings. Here few counterpossibilities are epistemically relevant and the ones depicted by wild skeptical scenarios are not even remotely envisaged. Further, for rudimentary conversational purposes, ‘I know that *p*’ may have the same force as ‘I have been informed that *p*’. Thus, the contextualist invites us to acknowledge how dysfunctional it would be in the routine inquiries of everyday life to put on the philosophical hat and consider people epistemically guilty until proven innocent. As a matter of fact, when not at work, philosophers have no trouble attuning ‘know’ to its pragmatic niche for attributing epistemic states to themselves and their fellows (even young children) to preserve and enhance effective interactions in the social world.

Now, the general worry. Why should an epistemological context be inevitably a *skeptical* one? *Within* the special context created ‘by doing epistemology’, the Lewisian contextualist ends up agreeing with the skeptic that the satisfaction of ‘know’ is outside the scope of the achievable. So, the skeptic instigating a suicidal behaviour of ‘knows’ and its cognates is conceded *carte blanche* precisely in the most conscientious investigation of *human* knowledge ascriptions. To preserve the parallel with ‘tall’, the contextualist looks like someone who holds that tallness ascriptions are context-sensitive, while constraining the sense of ‘tall’ to that used by NBA coaches when she is engaged in professional inquiry as to whether normal people are tall or not.

On the other hand, to the extent that Lewis’s view that ‘knows’ is abundant in a quotidian conversation context is still woven into the fabric of epistemology, and to the extent that it would be unprofessional of us to pretend otherwise, the question is how much elusiveness in this abundance we are left with. This question is independent from the familiar charge levelled against Lewisian contextualism: Lewisian ordinary conversations are vulnerable to the manipulation of a skeptical interlocutor, whose launch of a skeptical possibility of error is enough for knowledge to dissolve into air.<sup>6</sup> This charge is essentially correct but betrays an overreaction: arguably, the possibility that the skeptic crops up in a mundane conversational scene counts itself as a skeptical one. Rather we contend that in Lewis’s “story”—as he calls it—certain intrusive features of the evaluated/evaluator interaction imparts a neo-skeptical spin to knowledge ascriptions. By ‘neo-skeptical spin’, we mean two interlaced facts: (i) Lewisian evaluators *overgenerate* the possibilities of error which cannot be ignored; (ii) under disagreement with the evaluator about ignoring a given possibility, the evaluated is supposed to attend to it anyhow (on pain of not satisfying ‘know’). This ‘neo-skeptical spin’ produces a mismatch between our use of ‘know’ in ordinary life and the use of ‘know’ by ‘Lewisian’ ordinary speakers.

<sup>6</sup> Blome-Tillmann speaks for many in writing that for Lewis “any context in which one considers sceptical arguments is a context in which one does not properly ignore sceptical possibilities” (Blome-Tillmann 2009, p. 245).

A further and full-scale mismatch is closely tied to the claim that the negative capacity of pruning branches off the tree of alternatives cannot be the whole of the contextualist story about knowledge attribution. The value of getting *better* knowledge has gone missing. If knowledge admits of degrees—as Lewis claims—we would need to explain why better knowledge is better than knowledge. The task of the paper is that of addressing and solving, or substantially relieving these mismatches along the lines of a neo-Lewisian contextualist project. We argue that various threads should be untangled and gaps bridged in order for the Lewisian story to be applicable to the textures of human lives, if philosophy and ordinary life are to blend their concerns.<sup>7</sup>

Roadmap: in Sect. 2, we take a quick look at some basic tensions embedded in the evaluator’s application of the Lewisian rules which legislate on what is modally relevant or irrelevant to an epistemic appraisal. In particular, we argue that the magnitude of the work discharged by the *rise* of alternatives in Lewis’s theory turns out to be problematic in the absence of an explicit differentiation between ‘discovered’ and ‘invented’ alternatives, a differentiation which does have intuitive traction. In light of this distinction, we propose an amendment to the **Rule of Attention**, whereby relevance is not spread over all the alternatives we are attending to, but only over the discovered ones (**Rule of Discovery**).

In Sect. 3, we show that under the unrestricted **Rule of Attention** crucial moments of the interaction between two Lewisian interlocutors prove to be cognitively unrealistic. The subject is forced willy-nilly to divert her own attention to some alternatives proposed by the interlocutor. We argue that this ‘willy-nilly’ properly counts as an interference on the evaluator’s part, whilst it fails to explain the fact of life that when speakers disagree about the ‘ignoring of an alternative’, their conflict is precisely as to whether that alternative is *worth* being here and now attended to. In casual and concrete conversations possibilities are raised and eliminated, as well as *discounted* and *negotiated*. We argue that the **Rule of Discovery** promotes knowledge attribution as a *cooperative* enterprise, while explaining why a subject may attend to a possibility with one interlocutor, but not with another.

In Sect. 4, we dismiss the strange role assigned by Lewis to the elimination of *irrelevance* in improving knowledge. Then, we sketch a positive suggestion by taking into account another kind of everyday connection with modal states of affairs. Human beings are imaginers. In particular, we explore some reasons for introducing in the Lewisian original story a new vector of value—modal *extroversion*—which makes the satisfaction of ‘know better’ depend on the discovery of (actualized) possibili-

<sup>7</sup> Suppose a chasm between philosophy and ordinary life. Then, it would be perfectly obvious why skeptical possibilities are relevant *exclusively* in the special environment of the philosophy classroom, *if* they are relevant within it. But it would still not be obvious that they *are* relevant within it. Quite the contrary; the strong pull of skeptical possibilities, as an ‘intramural’ one, now cries out for explanation: why do they *interest* philosophers in the first place? Granted, under the chasm, subjects migration from the outside to the inside of the seminar would not affect their epistemic integrity, any more than our bankruptcy in the game of Monopoly affects our financial integrity. But then the very compromise between ‘elusiveness’ in philosophy and ‘abundance’ in daily life would hold vacuously. Perhaps it is unclear how to even make sense of the point of having the verb ‘know’ in philosophical parlance, if its usage is so idiosyncratic (unlike the verb ‘mortgage’, say, in Monopoly). In sum, it is clear enough that the contextualist has no choice but to conclude that philosophy is not divorced from everyday life.

ties, together with their elimination. Modal extroversion allows us to advance a clear explanation of why ‘know’ comes in degrees. In Sect. 5, we briefly take stock.

## 2 Lewis’s rules

A few reminders about the Lewisian rules that are lumped into two distinct packages: four prohibitive rules purport to capture the class of possibilities that the evaluator *cannot* properly ignore in a given context.<sup>8</sup> Three *permissive* rules fix the class of possibilities that *may* be sliced off in assessing whether someone knows or not. This is Lewis’s own order of presentation.

The **Rule of Actuality** explains why truth is not plugged into (Def<sub>c</sub>). The rule says that a  $\neg p$ -possibility that actually obtains is never properly ignored. Actuality, being a relevant uneliminated possibility, must be taken as *terra firma*.

The **Rule of Belief** forbids the evaluator to ignore any possibility that the knowing subject believes or ought to believe to obtain. The essential rationale for the rule is that we must not attribute knowledge that  $p$  to someone who, on the basis of her own beliefs, should have attended to a  $\neg p$ -possibility but refrained from doing so. Suppose we are reading a lipogram where every English word containing ‘e’ is eliminated. Informed of the fact that its author believes that the correct spelling of ‘whisky’ is ‘whiskey’, we should not properly ignore the absence of the word ‘whiskey’ in the text if we read a passage about liquors. Conversely, the presence of ‘whisky’ should be interpreted as a mistake.

This analogy makes visible that the **Rule of Belief** promotes a task which is in general impossible to live up to. Other selves do not look to us like umbrella-carriers on the street to whom we can attribute at least the belief that rain is possible. We seldom have access to information about undisclosed beliefs of other people, and the vast majority of these beliefs, arrived unreflectively, operates in an automatic way out of the sight of their owners. Such a fragmentary or inaccurate access to the doxastic life of others explains why two evaluators sharing the same conversational context and considering whether  $S$  knows that  $p$  might reach contradictory conclusions.

The **Rule of Resemblance** says that if a  $\neg p$ -possibility saliently resembles a possibility which is not properly ignored by the evaluator, then it should not be properly ignored either. This is the Lewisian antidote to the epistemic pollution which undermines knowledge ascriptions in Gettier cases and their variants. For instance, if an evaluator maintains that Bertie is the unwary victim of a Gettier scenario when he looks at a broken clock which happens to display the right time, then s/he cannot attribute to Bertie knowledge of the time: Bertie is indeed not properly ignoring the  $\neg p$ -possibility resembling actuality, according to which he is looking at a clock that stopped precisely twelve hours before.<sup>9</sup>

<sup>8</sup> Possibilities require some appropriate degree of specificity, not a ‘maximal’ one. What is ‘specific enough’, however, is vague matter. Lewis’s idea is that a specific enough possibility cannot be split into subcases dissimilar in some epistemologically relevant way.

<sup>9</sup> Cf. (Russell 1948, p. 170). The **Rule of Resemblance** admits an exception in order to avoid jumping into skepticism with both feet: skeptical possibilities cannot become relevant just because they resemble

Take also the familiar Aesop's fable: the shepherd boy eventually cries: "wolf! wolf!" when a wolf actually comes from the forest; but villagers do not stir to come to his aid, being fed up with his false alarms. Now imagine that a stranger, passing through the village, happens to hear the genuine cry "wolf! wolf!" of the shepherd boy. But the possibility resembling actuality that the current cry of the boy is a false alarm becomes salient to the villagers, and this is why they refrain from attributing to the stranger the knowledge that a wolf is about to attack the flock.<sup>10</sup>

Now, the package of permissive rules. The class of possibilities that may be sliced off contains: any possibility concerning breakdown of reliable processes (such as perception, memory, and testimony) (Rule of Reliability); any possibility about errors in non-deductive methods of inference such as sampling and abduction (Rule of Method); and any possibility generally ignored by people around us (Rule of Conservatism).

Finally, the Rule of Attention affirms that the evaluator may not properly ignore a  $\neg p$ -possibility to which *s/he* is attending.<sup>11</sup> This is far from being "more a triviality than a rule"—as Lewis says—because there is no intrinsic plausibility in the claim that relevance applies to *any* considered possibility just in virtue of being considered. The rule is "built to explain how the skeptic manages to sway us—why his argument seems irresistible, however temporarily" (Lewis 1996, p. 560), but it is intended to cover the irresistibility of any counterpossibility which triggers contextual shifts in every conversational context. Normally John may qualify as knowing that his old piggy bank contains no money, if it emits no clinking sound of coins when shaken. But John's knowledge becomes defeated when a conversational partner asks him: 'and what if *paper* money is in piggy bank?'. By the Rule of Attention, John is no longer properly ignoring this possibility.

## 2.1 *Intermezzo*: permissive rules in knowledge self-attribution

As the opposite of attending, Lewisian ignoring (L-ignoring) is a phenomenon of an *involuntary* kind: for a thinker to ignore something is to lack it in her current consciousness (Oakley 2001, p. 318). As many commentators have pointed out, L-ignoring does not exhaust the genus of ignoring: when people in Macau complain that local taxi drivers often ignore their raised hand, they intend to express the thought that taxi drivers are well aware of their raised hand, and it is exactly this awareness that allows them to ignore *deliberately* the signalling. Call this kind of ignoring 'disregarding-ignoring' (d-ignoring) (Oakley 2001, p. 318). Of course, what is entailed by reflection over a current application of a permissive rule involves an instance of d-

---

Footnote 9 continued

the possibilities that actually obtain by virtue of the fact that they cannot be eliminated by the subject's evidence.

<sup>10</sup> Thanks for an anonymous referee for pressing us to understand this example in relation to the Rule of Resemblance.

<sup>11</sup> "A possibility not ignored at all is *ipso facto* not properly ignored [...] No matter how far-fetched a certain possibility may be, no matter how properly we might have ignored it in some other context, if in this context we are not in fact ignoring it but attending to it, then for us now it is a relevant alternative" (Lewis 1996, p. 559).

ignoring. Whenever the Lewisian evaluator is mindful of the fact that she is following some rules determining the possibilities that are safe and convenient to ignore, she is thereby monitoring exactly those possibilities that count as negligible.<sup>12</sup>

It turns out, then, that L-ignoring cannot track the *usefulness* of the permissive rules with respect to ascriptions of knowledge to *oneself*. Speakers ordinarily employ the expression ‘I know that *p*’ when they think they are in possession of evidence for *p* that their audience lacks, or—as Malcolm would say in a Wittgensteinian vein—when other people have expressed doubts about *p* (cf. (Malcolm 1949)). But any Lewisian speaker claiming: ‘I know that it’s raining’ after consulting the Rule of Reliability would be automatically committed to *permissive d-ignoring*. One might thus suspect that the usefulness of the permissive rules in knowledge self-ascriptions does not lie in allowing the self-evaluator to L-ignore some possibilities, but rather in allowing her to *justify* her d-ignoring them, if pressed by someone to do so. And of course, here the Lewisian might feel uncomfortable, as space is opened up for some normatively loaded discourse. She can envisage for self-attributors something like an ‘off line’ application of the permissive rules. A meta-rule would have to state that self-attributors may properly L-ignore certain possibilities, provided that they also L-ignore the rules which license their L-ignoring of these possibilities as a proper L-ignoring. This is an awkward result, to say the least. In sum, what these comments invite the Lewisian to acknowledge is that knowledge self-attributions are *sui generis* knowledge attributions to the extent that any conscious application of some permissive rules boils down to a sort of justification for d-ignoring some possibilities.<sup>13</sup> So much for knowledge self-attributions.

## 2.2 ‘Discovered’ and ‘invented’ alternatives

Let us briefly discuss the division of labour between permissive and prohibitive rules. Lewis’s idea seems to be that the permissive ignoring carries an implicit claim of interpersonal authority, while constituting a sort of *ceteris paribus* ignoring: it remains in force until some prohibitive rules, especially the Rule of Attention, comes to the fore. Our contention is that the Lewisian story assigns this rule a too great workload. Attention is a psychological notion and, as noted above, the rule’s ambition is to make sense of the demands placed by the occurrence of the thought of a counterpossibility on our psychological economies. But without any qualification, the abrupt transition that is supposed to take place in the speaker’s mind from permissive to prohibited ignoring is not psychologically realistic. What does it mean that speakers in a given context become aware of a certain alternative to a contingent proposition? Lewis’s story is reticent on the process of ‘modal awareness’ formation. Yet this the axis on

<sup>12</sup> This point seems implicit in Lewis’s remark that “the epistemology we’ve just been doing, at any rate, soon became an investigation of the ignoring of possibilities. But to investigate the ignoring of them was ipso facto not to ignore them” (Lewis 1996, pp. 559–560).

<sup>13</sup> Lewis only admits that “the only place where belief and justification enter my story” is through the Rule of Belief (Lewis 1996, p. 556). The word itself ‘evidence’ seems to be a ghost of the link between justification and knowledge, which Lewis breaks without much ado. His quick argument amounts to the remark that since sometimes we don’t know *how* we know, we have in these cases *unjustified* beliefs that count as knowledge (p. 551). Yet, one might object that we could *unwittingly* have a justification for these beliefs. On this point see Neta (2003, pp. 11–12). More generally, one might also submit that ‘evaluation’ or ‘attribution’ connote a judgement with a belief-like structure.



which to consider whether alternatives are matter of *discovery*, or of *invention*. In other words, the **Rule of Attention** is oblivious to the distinction between ‘discovered’ and ‘invented’ alternatives.

The first thing to say is that there need not be anything elaborate in such distinction. One can glean a sound insight from examples. In the famous Dretske zebra case (Dretske 1970), we are evaluating whether Fred knows that the animal in the zoo he is looking at is a zebra. Yes, he does. Then, we are told that zoo authorities, to save money, have replaced some of the zebras with cleverly disguised mules. It seems fair to say that in this case we have *discovered* the alternative that Fred is not looking at a real zebra. Roughly speaking, the notion of ‘discovered alternative’ captures the intuition that alternatives are rarely raised in and for themselves but are instigated to advance our understanding of the world around us. The set of discovered alternatives is meant to contain alternatives both triggered by novel experiences and new information and licensed by memory states and special reasons.<sup>14</sup> By contrast, consider the case of Harold who is told by Anna that the temperature hovers around 30°C according to the weather report she has just heard on the radio. Harold protests, and invites Anna to consider the possibility that the weather report was delivered by a stranger who eluded surveillance, entered the radio studio, took hold of the microphone for a few minutes, and read *his own* report. In this case, we say that Harold has *invented* the alternative that the radio weather report is unreliable. What makes this situation implausible is that Harold presumes this possibility to be bothered about, yet he is entirely unperturbed by his inability to track the etiology of it. Alternatively, Harold would have to produce this etiology for his denial of knowledge to Anna to have any traction.

The issue here is really about the basic motor of the raising of alternatives in our cognitive processes. The unsophisticated claim we are pressing is just this: in the interim between ignoring properly an alternative and attending to it, an encounter with new information, or some handling of old information is required. We may *properly* ignore the alternative of seeing a mule painted to look like a zebra, inasmuch as normally Zoo visitors do the same. But it is precisely because such ignoring is *proper* ignoring that we shall not raise the alternative of seeing a painted mule instead of a zebra, without having *special* reasons for doing so. Now we should note that there is instead an intrinsic plausibility in the claim that relevance applies to any considered alternative in virtue of being *discovered*. That is to say, it is when we narrow down the range of alternatives to the discovered ones, that the **Rule of Attention** becomes “more a triviality than a rule”:

**Rule of Discovery:**

any discovered  $\neg p$ -possibility to which we are attending may not be properly ignored.

**Rule of Discovery** amounts to an anti-skeptical move, if one plausibly assumes that skeptical possibilities are invented possibilities rather than discovered possibilities.<sup>15</sup> This may offer comfort to those commentators who insist that the **Rule of Attention**

<sup>14</sup> For example, the possibility that the cake bought for the party is made without buckwheat flour counts as ‘discovered’ by Fred, if he remembers that one of his guests has a serious allergy to buckwheat.

<sup>15</sup> It would take us too far to pursue this point further, but the assumption may be defended on the basis of the evidential parity between skeptical possibilities and their counterpossibilities.

is skeptic-friendly in guaranteeing the success of skeptical improvisations within conversations. Moreover, the Rule of Discovery is conducive to cooperation between speakers. Or so we argue in the next section.

### 3 Gimpel the smart

The Rule of Attention, as it stands, is evaluator-sensitive to the point that the evaluator's persistence prevails over 'Humean' distractions:

If you bring some hitherto ignored possibility to our attention, then straightway we are not ignoring it at all, so *a fortiori* we are not properly ignoring it. How can this alteration of our conversational state be undone? If you are persistent, perhaps it cannot be undone - at least not so long as you are around. Even if we go off and play backgammon, and afterward start our conversation afresh, you might turn up and call our attention to it all over again (Lewis 1996, p. 560).

In the indeterminate space for its activity, however, the Rule of Attention may accommodate the whole gamut of effects from laughter to tears. Consider the case of Gimpel in Isaac Singer's short story *Gimpel the Fool*. One night, coming home unexpectedly, he discovers a man sleeping next to his wife Elka, but turns on his heels not to wake the child. When he goes to the rabbi for advice, he is ordered to divorce Elka at once. She, however, brazenly denies everything. Reluctant to divorce her, eventually Gimpel tells the rabbi that he must have hallucinated the scene. Now, let us modify Singer's story by imagining that Elka's husband is not Gimpel the Fool, but Gimpel the Smart. We have to consider two toy cases of conversational interaction, accordant and discordant:

- (Case 1) (Accordant). The rabbi mentions to Gimpel the alternative of hallucination, maybe to defuse the tension (and this intention is presumably reflected in the prosody of rabbi's utterance). Although this is not Lewis's case of a possibility cited *by mistake*, it is equally a case where—as he puts it—“we might quickly strike a tacit agreement to speak just as if we were ignoring [an alternative]; and after just a little of that, doubtless [that alternative] really would be ignored” (Lewis 1996, p. 560).
- (Case 2) (Discordant). Elka presents Gimpel with the alternative of hallucination to defend herself to the bitter end. Although there is no agreement between them to speak as if ignoring it, the alternative is *attended to* by Gimpel as well, according to Lewis. Rule of Attention should make it relevant. However, this alternative is strictly *irrelevant* to him. Our intuition here is that Gimpel can still correctly claim to know that his wife committed adultery.<sup>16</sup>

In either case, the broaching of the counterpossibility that Gimpel the Smart might have experienced an astonishingly clear hallucination does not make him retract his

<sup>16</sup> A similar case has been discussed in (Blome-Tillmann 2009, p. 246) and (Kenna 2009, p. 743). Kenna understands it as showing that “an alternative that is being attended to in a context [...] can be irrelevant” (*ibidem*). Our reading is slightly different (see *infra*).

earlier claim that he knows that his wife has been unfaithful. Nevertheless, in **Case 1** the counterpossibility is not brought up as a challenge to Gimpel's knowledge claim. It has a distinctly ironic effect, as the rabbi does not question Gimpel's version of the facts. Soon, Gimpel and the rabbi will silently (and accordantly) let go of such counterpossibility. A case of this type illustrates that Lewis is not committed to the idea that an error possibility—and a skeptical one in particular—*must* be taken into account every time an interlocutor conjures it up.<sup>17</sup> The moral here is that in a dialogical exchange, a possibility may wax and wane along with empathy. More generally, our modal talk is affected by a large array of purposes in human communication: counterpossibilities are also brought up to match the others' affective states, under the guise of compassion, irony, humour, mockery, and sarcasm, so that they work for the speakers as an effective tool for emotional attunement without alethic commitment. Accordingly, any attempt to regiment the subtleties of our everyday uses of modal talk comes with a ziggurat of hard policing assumptions of a descriptive and normative kind. Although the contextualist should acknowledge this 'pragmatic ground noise', it doesn't fall into the proper domain of epistemology.

**Case 2** is a counterexample to the **Rule of Attention** by exploiting a puzzling Lewisian condition, which we may call the *willy-nilly condition* (**WN**, for short). Let  $P$  the proponent of an alternative, and  $D$  the dissenter. Then, for any alternative  $w$ :

(**WN**) Under disagreement with  $P$  about the ignoring of  $w$ ,  $D$  is supposed to attend to  $w$  anyhow.

Conjoined with the **Rule of Attention**, (**WN**) implies willy-nilly relevance:

(**R<sub>WN</sub>**) Under disagreement with  $P$  about the ignoring of  $w$ ,  $w$  is supposed to be relevant to  $D$ .

In **Case 2**, Gimpel satisfies (**WN**) by considering the possibility of hallucination. Yet, he does not take it to be a relevant one, against (**R<sub>WN</sub>**).

Lewis offers no pragmatic justification for (**WN**) beyond the appeal to the needs of cooperative conversation (Lewis 1996, p. 560). Nevertheless, the sort of cooperation he has in mind exhibits lack of reciprocation. There is no simple argument, indeed, for the conclusion canvassed by Lewis that the effort of cooperation is only required on the side of the interlocutor who is asked to consider a certain formerly ignored possibility. Why should it be *fair* to confront a possibility anyway in the face of an *actual* disagreement? This one-way traffic cooperation, we submit, is an unrealistic result.

Cooperation in speech dynamics has a protean nature whose complexity escapes a Gricean lens. For example, normally people are charitable in judging what their interlocutors mean, but they do tend to be much less charitable in judging whether their interlocutors are *right*. And yet, refraining from outspoken dissent is an underappreciated form of conversational cooperation. People also tend to be quiet when their conversational partners all of a sudden change the topic of conversation or repeat what already said. Such a low-degree cooperation descends not so much from fair play as from the awareness that the course of a typical conversation is something

<sup>17</sup> This has been noted by Kenna (2009, p. 743).

unpredictable, fraught with contingencies, vagaries, and redundancies. Nonetheless, it is one thing to consent to other people's deviations from the conversational path or repetitions, and quite another to self-impose the special attention required for taking into account an alien possibility, which is perceived as utterly forced with respect to that domain of discourse and to that interlocutor. Actually, this effort is unusual, being closer to a capitulation to the other party's modal scruples than to a cooperative concession. For instance, when we are pressed to consider the possibility that we are looking at artificial flowers, we feel entitled to discount such possibility if we know that the owner of the house detests artificial flowers. Technically, our evidence does not eliminate this possibility, whose obtaining is still compatible with the information about the owner. What our evidence mandates is the *rejection* of this possibility as a good candidate for eliminability.

In sum, if speakers find themselves in genuine disagreement about 'ignoring an alternative', it is because they hold different doxastic attitudes towards the legitimacy of embarking on a certain mental *action* here and now. Like Macau taxi drivers's d-ignoring, people's d-ignoring allows a *decision* to be made about whether the action of attending to that possibility is worth doing.<sup>18</sup>

At any rate, it is with the help of (WN) that the Rule of Attention propels towards skepticism: the skeptic can prevail just by coming on the scene and playing his cards, even if we do not want to play. From the premise that we do not want to play the skeptical game, by (WN) the most that we can conclude is that we play reluctantly. In claiming that (WN) is puzzling, however, we mean that it is so independently of any skeptical inclination of the speaker. And, as noted, the possibility for us to come across a skeptical interlocutor in an ordinary conversation is itself a skeptical possibility. Ordinary life encounters with disagreement about the ignoring of an alternative typically concern cases where one proponent is supposed to lack open-mindedness, freedom from bias, or sincerity in raising that alternative. In a wide range of cases one's ostrich-like attitude may manifest itself in *proposing* a possibility rather than discounting it.<sup>19</sup> In real-life situations, in other words, we suppose that what makes it unreasonable for one side to retain such a defensive attitude is precisely what makes it reasonable for the other side to choose *not* to fight that modal battle. If these remarks are correct, then we see why it would be unnecessarily heroic for Gimpel the Smart, in *talking with Elka*, to entertain the possibility of having suffered a hallucination on that night.

On the other hand, it is entailed by the description of Case 2 that the disagreement between Elka and Gimpel is not a genuine one. There is no conflict between them that is both interesting and epistemic. In her heart, she *attributes* to the husband knowledge that she committed adultery. Case 2, then, might suggest an amendment to the Rule of Attention by the following corrective: any counterpossibility to which

<sup>18</sup> Here Lewis misconstrues the role of d-ignoring, which he calls "make-believe ignoring": "we would ignore the far-fetched possibility if we could – but can we? Perhaps at first our attempted ignoring would be make-believe ignoring, or self-deceptive ignoring; later, perhaps, it might ripen into genuine ignoring. But in the meantime, do we know? There may be no definite answer..." (Lewis 1996, p. 560).

<sup>19</sup> To take a dramatic example, consider how in the 1950s the Tobacco Manufacturers Standing Committee used R. A. Fischer's claim that scientists did not know that smoking causes cancer since they did not eliminate the possibility that a smoker's gene causes *both* a desire to smoke and lung cancer. See Sorensen (2004).

we are *sincerely* attending should not be properly ignored. Yet, this is an anti-Elka amendment but not an anti-skeptical one, inasmuch as the construal of the skeptic as an insincere interlocutor is not mandatory. The skeptic can be considered to be sincere, and still he is modally indiscrete.

However, different responses to the same alternative can be appropriate, when the alternative is raised by different speakers with different intentions. In short, *D* may not satisfy (WN) with *S*, while satisfying it with *S'*. Consider this case:

(Case 3) When Gimpel goes to the rabbi for advice, he senses that Elka said untruly to the rabbi that Gimpel suffers from hallucinations. There is no agreement between Gimpel and the rabbi to speak as if ignoring the possibility of hallucination, but the alternative is *attended to* by Gimpel anyway.

Why? Because he thinks that the rabbi, unlike Elka, has (weak) reasons to think that the possibility of hallucination obtains, and so Gimpel thinks that the rabbi has reasons not to attribute to him knowledge that Elka committed adultery. In short, that alternative enters the scene as discovered by the rabbi, unlike the one invented by Elka, and Gimpel recognizes the difference. Of course, he has no intention to re-evaluate what happened that night, but seeks an agreement about *d*-ignoring with a person very influential on the community: relevance may be interlocutor-sensitive.

This means that invented and discovered possibilities may have the same content, while playing different roles for a dissenter grappling with different counterparts. So, the heart of diagnosis is this. Case 2 shows that an *invented* alternative that is being attended to in a context can be irrelevant. Similar cases can be circumvented by the Rule of Discovery.

Accepting this rule commits us to (WN<sup>d</sup>):

(WN<sup>d</sup>) Under disagreement with *P* about the ignoring of the discovered alternative *w*, *D* is supposed to attend to *w* anyhow.

(WN<sup>d</sup>) is realistic and familiar. In particular, like Gimpel the Smart in Case 3, a dissenter satisfies (WN<sup>d</sup>) when she seeks an argumentative contact with the proponent. The Lewisian will find in it grist for the normative mill. Yet real-life dialogues unroll far beyond the frugal conversation which for Lewis exactly serves to prevent speakers from entering into argumentative contact. When conversation brings to light discovered possibilities, frugality is unrealistic, because discovered possibilities have the psychological punch that the invented ones lack.

## 4 Imagining possibilities

### 4.1 Prized irrelevance

*Pace* Dretske—who holds that propositional knowledge is an all-or-nothing affair like being pregnant and unlike being wealthy (Dretske 1981, p. 363)—the Lewisian contextualist is inclined to think that knowledge-that comes in degrees, so that the analogy between ‘knows’ and gradable adjectives like ‘rich’ resurfaces.<sup>20</sup> Lewis’s

<sup>20</sup> However, the analogy between ‘knows’ and context-sensitive adjectives, like every analogy, sooner or later breaks down: when people disagree about whether the predicate ‘knows’ is instantiated in a given

own example involves himself searching for Possum the cat. If David looks all around the room, then he may be described as knowing that Possum is not in the room, even if he doesn't check the drawers of the desk; the possibility that Possum jumped into a drawer may be properly ignored, indeed. Nonetheless, if David also checks all the desk drawers, then he knows *better* that Possum is not in the study. In this second case, David's enhancement of his own knowledge rests "more on the elimination of not-P possibilities, less on the ignoring of them" (Lewis 1996, p. 562).

On the other hand, the elimination of new *relevant* alternatives fatally makes the subject's epistemic state just 'move' to a new context, instead of improving it.<sup>21</sup> To sum up: in a given context, the more *irrelevant*  $\neg p$ -possibilities are eliminated by our evidence, the better our knowledge that  $p$ . If this is the insight, where does it lead us? The alleged importance of irrelevance should give us pause: why should the elimination of an irrelevant possibility have anything do with enhancing knowledge? Our primary reason for dissatisfaction with this view is that it is extremely obscure. The Lewisian is charged with the burden of sorting out how, and not merely assume that, a certain eliminated possibility may be a condition of epistemic progress, while continuing to be irrelevant.

It emerges soon that what Lewis wants us to take seriously is the property of *stability*: "better knowledge is more stable knowledge: it stands more chance of surviving a shift of attention in which we begin to attend to some of the possibilities formerly ignored" (Lewis 1996, pp. 562–563). This is not a tribute to Plato's *Meno*, however, where Socrates invokes stability to explain why knowledge is more valuable than doxastic states which fall short of knowledge. Lewis appeals to stability in order to show why to know better that  $p$  is more valuable than to know that  $p$ . Nonetheless, again, it is unclear how to spell out the idea that the elimination of irrelevant possibilities is conducive to stability. One might try to defend this idea as follows: in a mundane shift of attention,  $S$ 's evidence has more chance of ruling out some "possibilities formerly ignored" if such evidence has been *enlarged* by previous elimination of some *irrelevant* possibilities, i.e. possibilities that  $S$  was entitled to discount; this is why the elimination of irrelevance is active in leaving  $S$ 's knowledge unscathed after the shift. For instance, if the possibility that Possum has been locked into a drawer by some malicious colleague of David's becomes relevant, then David can eliminate it on the spot, since he has already checked all the drawers.

However, this suggestion does not withstand scrutiny. First, the above sketched cognitive trajectory is severely affected by redundancy. If David has *already* inspected the drawers when the possibility that Possum was inside was irrelevant, why should

---

Footnote 20 Continued

circumstance, they consider the dispute as in principle amenable to resolution. But in case of dissonant beliefs about the correct application of 'elegant' disputants do not presume to resolve the issue.

<sup>21</sup> Consider a case involving a shift in context. Dancer Olga lives in Vologda, in the Northwest of Russia; the next day she is going to an audition for the Bolshoi Ballett in Moscow. After looking at the train timetable on the railway website, she may truly claim that she knows that  $p$  [= the train to Moscow leaves at six a.m.]. Suppose that Olga raises the possibility that the railway website is unreliable. By raising this possibility in her soliloquy, Olga *changes* the context to a new one in which this uneliminated possibility is relevant. Thus, she doesn't know better, as she has simply changed the context. For a discussion of comparative knowledge in Lewisian contextualism, see Douven (2004).

he “begin to attend to” the possibility that Possum has been locked into a drawer by someone? Second, if the formerly eliminated irrelevant possibility acquires new currency, that is fortuitous for it. And the obvious implication of this fortuitousness is that the concomitant stabilization of knowledge will come to fruition as an unpredictable happening. Hence, more work needs to be done to vindicate the claim that the elimination of irrelevant possibilities is conducive to stability.

Cursorily understood, at any rate, the function of eliminating irrelevant possibilities might rather seem to be that of *mitigating* the impact of attention shifts on our temporary knowledge: the more irrelevant  $\neg p$ -possibilities our evidence eliminates, the less our knowledge that  $p$  succumbs to subsequent shifts of attention. On the whole, anyway, a function of this sort constitutes a *protection* for knowledge instead of an amelioration of it. Lewisian stability is entirely a defensive property. So, why is stability valuable? Lewis’s answer is: because it abates knowledge’s fallibility, or precariousness. Socrates’s answer is: because it generates a tether, i.e. *connectivity*. Socrates’s answer implies that of Lewis: fallibility stands in inverse proportion to the ramifications of knowledge connections. Now let us imagine for the sake of argument that an infallible oracle would declare that Socrates’s answer is the correct one, but that we are required to provide a plausible articulation of the tether. At the moment, the oracle’s verdict would be uncongenial to the contextualist. It seems obvious that the tether cannot be achieved by sifting through eliminated possibilities, be they relevant or irrelevant. Why? Here is a two-step argument: (1) eliminated counterpossibilities are unactualized possibilities; therefore, (2) unactualized possibilities, as spatiotemporally disconnected from actuality, cannot be a point of entry into a network of positive connections. We think, however, that contextualism possesses inner resources to meet the challenge posed by the oracle, and to make sense of the thesis that knowledge comes in degrees. This is what the next two subsections of the paper aim to briefly show.

## 4.2 The contextual efficacy of stupidity

Some possibilities ignored by the epistemically lazy person will not be properly ignored. Well and granted. But what if knowledge is somehow constrained by laziness itself, i.e. it turns out to be a symptom of the lack of intellectual passions? In a paper where she expresses her own disenchantment with the rivalry between internalism and externalism in mainstream epistemology, Catherine Elgin invites us to consider a case like the following (Elgin 1988, pp. 304–305).

(W&H) The amateur ornithologist John H. Watson and the expert one Sherlock Holmes are both looking at a strange bird on the window ledge: dark brown wings, a pale pink plumage on its head, shoulders and underside, along with pink feet and beak. Unbeknownst to them, that bird has just escaped from the London Zoo. Watson unhesitatingly recognizes it to be a pink pigeon. He ignores the bird’s habitat, true. But the visual impression caused by the bird perfectly matches with the picture on the manual of ornithology, which he is used to leaf through, skipping learned captions. Therefore, he confidently claims that he knows that  $p$  [= the bird on the ledge is a pink pigeon]. Now

Holmes. To him too the bird looks terribly like a specimen of pink pigeon. Yet, among the many things he knows of pink pigeon, there is the plain fact that it is a rare bird endemic to Mauritius, more than 6000 miles from London. Therefore, Holmes with admirable modesty admits that he doesn't know whether *p*. (*to be continued*).

The contextualist's reaction to (W&H) is that, in effect, Watson should be credited with the *knowledge* that *p*, if all one cares about is elimination of possibilities. In fact, one may assume, there is no counterpossibility that Watson actually believes to obtain that for an evaluator would be a violation of the Rule of Belief to ignore. Watson's evidence suffices to eliminate the possibility that the bird might be any other bird whose image is faithfully depicted in the manual, and these are the only relevant counterpossibilities. But, then, the moral to draw from (W&H) is a discouraging one, since, admittedly, there is something inadequate about Watson's position towards the proposition in question. Not despite, but *because* of his lack of intellectual curiosity, Watson is in the position to know what the expert Holmes ignores, while being exposed to the same body of perceptual evidence. The 'epistemic efficacy of stupidity'—as Elgin calls it (Elgin 1988)—forms something of an embarrassment for any view of knowledge attribution. Accordingly, this embarrassment is not special to the contextualist view, but attaches to it as well.

Now it is open to the contextualist to offer a more fine-grained reading of (W&H). Probably, s/he would suggest, it is as if 'Watson knows that *p*' and 'Holmes does not know whether *p*' were in two different contexts of knowledge attribution governed by low and high ornithological complexity, respectively. Namely, it is as though Watson, without worrying too much about the costs of being wrong, occupied a conversational setting with his young and curious nephew asking: 'What bird is that?', whereas Holmes is in a demanding context shared with other punctilious experts dwelling on the possibility of error against a set of stringent ornithological considerations. Accordingly, the lack of intellectual curiosity grants Watson no advantage over Holmes because they do not contend for the same prize. Our reply is that such response sounds too ecumenic by failing to do justice to what the example is drawing our attention to. It may be true that Watson and Holmes do not contend for the same prize: Holmes has more than Watson at stake as to whether *p*. But it would be a mistake to think that they are playing the same *eliminative* modal game. This immediately prompts the following question: can one make sense of the difference between the two games while still remaining within the bounds of contextualism?

It seems that we can. Contrary to the victim in a paradigmatic Gettier case, Watson may not be said here to be *lucky* that his belief that *p* turns out to be true. Indeed, the explanation for why he holds the belief that *p* entirely *depends* on what makes the belief that *p* true. Which is to say that his bird identification skill is far from being stochastic, but rather he is somewhat lucky that, in Rescher's phrase, his ornithological ignorance has its compensation (Rescher 1999). And how about the compensation of Holmes's background knowledge? The sequel of (W&H) is (W&H)<sup>+</sup>.

(W&H)<sup>+</sup> Yet Holmes does "not rank modesty among the virtues" (from *The Greek Interpreter* by Arthur Conan Doyle). His ruminations on pink pigeon habitat eventually push him towards the possibility *w* that the bird may have



escaped from the local zoo. On the basis of his sensory inputs and memory states and his belief that  $w$  is actualized, he concludes that he knows that  $p$ . (*End*).

Extracting a moral from (W&H) and (W&H)<sup>+</sup>, we may say that the central difference between Watson's and Holmes's achievements pertains not so much to their negative capacity of eliminating possibilities, as to their positive capacity of *imagining* or *generating* possibilities. Holmes's proper grasp of the actual-world structure is a fertile ground against which to imagine the possibility  $w$  that the bird may have escaped from the Zoo. The word 'end' in (W&H)<sup>+</sup> is intended to signal that Watson and Holmes's epistemic positions cannot be equalized, supposing that Holmes simply presents  $w$  to Watson. Here the Rule of Attention and the Rule of Discovery remain inert:  $w$  is compatible with Watson's knowledge that  $p$ , whereas  $w$  overturns the reason why Holmes does *not* endorse  $p$ . Or, to put it in another way,  $w$  is irrelevant to Watson by virtue of his ignorance about pink pigeons, whereas  $w$  is operative in regulating Holmes's doxastic performance relative to  $p$ . But the basic point is just that Watson misses out on a sense of amazement concerning what he has seen: he does not have access to the conceptual resources that would allow him to imagine  $w$ , or appreciate  $w$  when presented by someone else.

In sum, Holmes knows that  $p$  *better* than Watson, and (Def<sub>c</sub>) frustrates his epistemic superiority. Lewisian contextualism lacks an account for the difference between Watson's and Holmes's epistemic states.

### 4.3 A new vector of value

To make matters more precise, let us fix some terminology. Let  $S$  be an agent that believes that  $p$ , or wonders whether  $p$ . We call a possibility  $w$  *abductively connected* with  $p$ , if: (i)  $w$  actually obtains; (ii)  $S$  thinks  $w$  as actualized; and (iii)  $S$ 's doxastic attitude towards  $p$  is strengthened by (ii).

Conditions (i) in concert with (ii) just affirms that  $S$ 's imagination cannot be a flight from the actual world: the content of  $w$  needs to be anchored in actuality. Condition (ii) automatically makes the possibility  $w$  relevant to whether  $p$ , insofar as  $w$  is linked with  $S$ 's doxastic attitude towards  $p$ . For example, if Mary Morstan thinks as actual the possibility that the pink pigeon has been placed in the London sky by a demon-helicopterist who wants to trick the ornithologists of the city, that possibility would satisfy condition (ii) and (iii), but not (i). On the other hand, the possibility that the pink bird on the window ledge is a common pigeon fed upon food containing high rates of beta carotene and canthaxanthin would not meet any of the conditions. Rather it is an instance of possibility that is spotlighted by the Rule of Attention or by the Rule of Discovery.<sup>22</sup> However, in another scenario, this very possibility counts as abductively connected with the proposition believed by Mary that the pink bird is a common pigeon, if she rightly thinks that this possibility is actualized.

<sup>22</sup> Readers of the online article 'Pink pigeon in London baffles bird experts' ([www.telegraph.co.uk](http://www.telegraph.co.uk), 10 August 2012) would arrive at that possibility through the Rule of Discovery.

Once the notion of abductive connection is taken on board, space is cleared for the introduction of two complementary properties relative to someone's evidence. Let us coin the terms 'modally *extrovert*' and 'modally *introvert*' for them. We say that *S*'s evidence is modally *extrovert* with respect to a certain proposition *p*, if *S*'s evidence generates possibilities abductively connected with *p*. Otherwise, *S*'s evidence is modally *introvert*. To 'generate' here should be understood as the opposite of 'eliminate'. Note that from the above definition it immediately follows that every possibility generated by extrovert evidence is a discovered one.<sup>23</sup> In light of this discussion, the next step is to incorporate the notion of belief into the definition of the satisfaction of 'knows'.

(Def<sub>c+</sub>) *S* satisfies 'knows that *p*' in context *c*, iff, for some *S*'s evidence *E*:

- (i) *S* believes that *p* on the basis of *E*;
- (ii) every  $\neg p$ -possibility is either eliminated by *E*, or properly ignored in *c*.<sup>24</sup>

Lewis offers an inconclusive argument to the effect that knowledge does not entail belief.<sup>25</sup> At least part of the explanation for the absence of belief in (Def<sub>c</sub>) may have to do with the fact that any uneliminated  $\neg p$ -possibility tends also to be a mark of the loss of the very belief that *p*. If Fred retracts his own knowledge that the animal in the pen is a zebra, as he cannot eliminate the possibility that it is a cleverly disguised mule, then Fred also recedes from his own *belief* that the animal in the pen is a zebra. But when one has to explain what makes people's knowledge become *better*, it is not worth the candle to play without the notion of belief: people's belief states are like orbital stations from which their modal investigations depart. To explain why Holmes knows that *p* better than Watson does, it is hard not to suppose that what is modally insulated is Watson's *belief state* with the content *p*, although this state is formed in a rational way. In our contextualist story, however, it suffices to assume that beliefs are means by which information is stored: information may be defended via the elimination of alternatives just as it may be connected to regions of modal space. This connection makes the difference as to what cognitive achievement one ends up with.

Now, it comes as no surprise that Holmes knows that *p* better than Watson, inasmuch as both satisfy 'knows that *p*' in (Def<sub>c+</sub>), but only Holmes satisfies this further condition:

- (iii) *E* is modally extrovert with respect to *p*

Let us turn to Lewis's own toy example: given the same context, it is not when David's evidence eliminates irrelevant alternatives in addition to the relevant ones that he knows better than his former self (or than another person) that Possum is not in the study.

<sup>23</sup> On a recent discussion of the relationship between abduction and modality see Biggs (2011).

<sup>24</sup> As concerns the rules fixing the membership to the set of relevant possibilities, Rule of Discovery replaces Rule of Attention.

<sup>25</sup> Lewis quotes approvingly Colin Radford's case of "the timid student about *p* who knows the answer but has no confidence that he has it right, and so does not believe what he knows". This is striking: the timid student's evidence, by hypothesis, is powerless to exclude the possibility that  $\neg p$ , otherwise the student would not be timid or unconfident at all. See also Ichikawa (2011) on putting a belief condition in a neo-Lewisian contextualism.

Rather, he knows better that Possum is not in the study when his evidence *generates* the actual possibility, specific enough, that Possum is elsewhere.

Modal extroversion, of course, admits of degrees. Reasonably enough, the degrees of modal extroversion vary in correspondence to the size of the set of the abductively connected possibilities that are generated by the subject's evidence. In sum, the elimination of possibilities seems to fall short in determining whether the predicate 'knows' is gradable: elimination is not enough, but one has to set the bar a good bit higher.

## 5 Concluding comments

'Elusive Knowledge' (1996) has reached the age of nineteen. This paper comes across in the Lewisian atelier as an intricate painting, which still requires studio assistants to take care of the cluster of details that make up the whole. These details do not merely appear as background battle scenes, for they sometimes prove to be crucial in orienting our vision along several dimensions. Yet the painting remains captivating.

Lewis's basic idea is that the evaluator behaves more as one participant to, than as one spectator of, the epistemic dramas of other people. This contrasts with the 'Panopticon' conception of knowledge attribution, according to which the evaluator is an inspector observing the evaluated from a vantage-point that remains unseen.<sup>26</sup> We believe that Lewis's thought is essentially correct, but some parts of its implementation are problematic. Our first warning has been that the very participation of the evaluator, on Lewis's view, may escalate into interference. This interference takes the form of an overgeneration of possibilities which are incumbent upon the evaluated. To narrow much of the distance that separates our use of 'know' in daily life and that of 'Lewisian' ordinary speakers we have proposed to replace the unrestricted *Rule of Attention* with the *Rule of Discovery* selecting discovered possibilities.

Moreover, the Lewisian original story is exclusively a negative one. We have instead investigated the prospects for a 'positive' account of contextualism, whereby knowledge attribution is not exhausted by the subject's negative capacity of eliminating the right range of possibilities, but it is also sensitive to her modal imagination. Lewis's elimination of irrelevant alternatives, indeed, seems to be explanatorily impotent with respect to the degrees of knowledge. The activity of imagining possibilities anchored in actual experience is one that is routinely exploited in everyday life. In our story, what we have called modal extroversion serves as the criterion for exploring the extent to which the knowledge predicate itself exhibits the property of being gradable: actuality is not only a relevant possibility which stays uneliminated, but it is also a relevant uneliminated possibility awakening, as a matter of imagination, other realized possibilities, which are connected into a coherent whole. This is the tether. This is what makes knowledge stable, whilst it remains contextual.

<sup>26</sup> What characterizes the post-Gettier discussion of a barrage of epistemic luck attributions vs. knowledge attributions is an information asymmetry: crucial information about a sequence of events, which pertains to the subject's environment, is available to the evaluator while being withheld from the subject. This information asymmetry is nicely discussed in Fogelin (1994).

## References

- Biggs, S. (2011). Abduction and modality. *Philosophy and Phenomenological Research*, 83, 283–326.
- Blome-Tillmann, M. (2009). Knowledge and presuppositions. *Mind*, 118, 241–294.
- Cohen, S. (1987). Knowledge, context, and social standards. *Synthese*, 73, 3–26.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52, 913–929.
- Descartes, R. (1984) *The Philosophical Writings of Descartes*. (J. Cottingham, R. Stoothoff, and D. Murdoch, Vol. II., Trans.). Cambridge: Cambridge University Press.
- Dodd, D. (2010). Confusion about concessive knowledge attributions. *Synthese*, 172, 381–396.
- Douven, I. (2004). The context-insensitivity of ‘Knowing More’ and ‘Knowing Better’. *Canadian Journal of Philosophy*, 34, 313–326.
- Dretske, F. I. (1970). Epistemic operators. *Journal of Philosophy*, 67, 1015–1016.
- Dretske, F. I. (1981). The pragmatic dimension of knowledge. *Philosophical Studies*, 40, 363–378.
- Elgin, C. (1988). The epistemic efficacy of stupidity. *Synthese*, 74, 297–311.
- Fogelin, R. J. (1994). *Pyrrhonian reflections on knowledge and justification*. New York: Oxford University Press.
- Ichikawa, J. (2011). Quantifiers and epistemic contextualism. *Philosophical Studies*, 155, 383–398.
- Kenna, R. (2009). Interests contextualism. *Philosophia*, 9, 741–750.
- Lewis, D. (1986). *On the plurality of worlds*. Oxford: Blackwell.
- Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, 8, 339–359.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–567.
- Malcolm, N. (1949). Defending common sense. *Philosophical Review*, 58, 201–220.
- Neta, R. (2003). Contextualism and the problem of the external world. *Philosophy and Phenomenological Research*, 66(1), 1–31.
- Oakley, I. T. (2001). A skeptic’s reply to Lewisian contextualism. *Canadian Journal of Philosophy*, 31, 309–332.
- Rescher, N. (1999). On learned ignorance. *The Southern Journal of Philosophy*, 37, 479–493.
- Russell, B. (1948). *Human knowledge: Its scope and its limits*. London: Allen and Unwin.
- Sorensen, R. (2004). Commercial applications of skepticism. In Walter Sinnott-Armstrong (Ed.), *Pyrrhonian skepticism*. Oxford: Oxford University Press.
- Unger, P. (1975). *Ignorance. The case for scepticism*. Oxford: Clarendon Press.