Prolegomena 12 (2) 2013: 381–412

Knowledge Attributions and the Psychology of Reasoning: A Case against Contextualism

ZVONIMIR ČULJAK / DRAGANA SEKULIĆ

Department of Philosophy, Centre for Croatian Studies of the University of Zagreb
Borongajska 83d, Zagreb, Croatia
zculjak@hrstud.hr / dsekulic1@vip.hr

ORIGINAL SCIENTIFIC ARTICLE / RECEIVED: 21–05–13 ACCEPTED: 09–11–13

ABSTRACT: Epistemic contextualism in the works of S. Cohen, K. DeRose, D. Lewis and others amounts to the *semantic* thesis that the truth conditions of knowledge attributions or denials vary according to the contextually shifting standards for knowledge attributions and to the indexical character of the predicate "knows". This semantic variation is primarily due to the *pragmatic* features of the attributor context, depending on "what is at stake" for the attributor. In this paper contextualism is confronted with some *invariantist* objections. These objections are supported, first, by the considerations of the alleged, but indeed not purely the semantic or meta-linguistic character of the main contextualist theses: it is argued that contextualism unavoidably descend to the object level, making certain substantive claims about knowledge, and that the ambiguous evidence of contextualist thought-experiments make the truth-oriented or intellectualist invariantist alternative a more plausible and more coherent view. Secondly, they are supported by the well-known empirical evidence from *cognitive psychology* by e.g. P. Wason, P. Johnson-Laird, D. Kahneman, P. Slovic, A. Tversky as well as by P. Cheng, J. Holland, K. Holyoak, R. Nisbett and P. Thagard. Those findings demonstrated that in various contexts people mostly do not reason according to logical or probabilistic rules, but according to some contextually convenient reasoning patterns, the reliance on which may lead to systematic logical or probabilistic errors. Their inferential performance has been obviously assessed according to the logical or probabilistic rules as *invariant standards* for the *attribution* of logical or probabilistic knowledge. Accordingly, it is argued in this paper that the change of truth conditions and truth values of respective inferential knowledge attributions or denials is sensitive to the changing facts in the subject context, and may not be explained by the shifting standards for knowledge attributions. These *standards* remain the same across the contexts even when practices in different contexts in fact follow some other reasoning rules (e.g. pragmatic reasoning schemas or heuristics). So, the varying truth conditions of the knowledge attributions in such cases depend only on the variations in the subject context.

KEY WORDS: Contextualism, inferential knowledge, invariantism, knowledge attributions, psychology of reasoning.

I. Introduction: contextualism vs. invariantism

The proponents of the so-called *attributor contextualism* in epistemology (hereafter: contextualism) like S. Cohen (1998, 1999, 2005), K. DeRose (1992, 1995, 2002, 2005, 2009), M. Heller (1999), D. Lewis (1996), and others, maintain that the contents and the truth-values of knowledge claims, knowledge ascribing or knowledge denying sentences, vary contextually in dependence of the *contextually shifting standards* for knowledge attributions. According to contextualism, the truth conditions of knowledge claims are sensitive to the context of the knowledge *attributor*, i.e. to the context in which knowledge claims are *uttered*. So, the truth-conditions of knowledge claims depend on *standards* or *facts* in the context in which these *attributions* are made. According to some contextualists, for example, according to S. Cohen (2005), it implies also that whether the knowledge claims "S knows that *p*" or "S does not know that *p*" are *true* or *not* depends on the features of the attributor context. So, whether a subject has a strong or a weak *epistemic position* toward *p* depends on whether high or low standards for epistemic claims are applied. But for contextualists that does not mean that these standards are invariantly epistemic in character or that they are straightforward knowledge standards: they function primarily as the *conversational* norms for the correct use of the predicate "knows" in respective conversational contexts, the contexts in which epistemic claims are made. "Knows" appears to be context sensitive like *indexical* terms "I" or "now". Applying D. Kaplan's (1989a, b) distinction concerning indexicals, contextualists maintain that the *character* or *linguistic meaning* of "knows" stays the same, but its *content* varies contextually. Consequently, when their context changes, knowledge claims change their contents, which are their respective propositions.

Some contextualists, DeRose and Cohen notably, claim also that the contextual variations of the contextual standards, which make "low standard" or "high standard" cases, depend on or are triggered by "what is at stake" in those contexts. It means that these standards are at the same time the functions of the salient *pragmatic*, non-objective and interest-related features of the context in which epistemic claims are uttered. According to Cohen (1999: 57) these pragmatic features include "the purposes, intentions, expectations, presuppositions, etc., of the speakers who utter these sentences". Contextualists find the supportive empirical evidence in the *linguistic data* provided by the *contextual semantics* of the verb "know" as well as in *pragmatic* assessments (in fictitious cases) of the believer's epistemic position and the consequences of her beliefs in respect to the overall practical situation in which the epistemic *attributor* is situated.

For the purposes of this paper we will summarize and discuss contextualism as a conjunction of the following two theses:

(i) the truth conditions and the truth values of knowledge claims as well as the semantic content of epistemic predicates like "knows" are sensitive and variable according to the context of the knowledge *attributor* (hereafter: *the semantic thesis*), and

(ii) these semantic variations depend on certain *pragmatic* features, i.e. on "what is at stake" in the attributor's context (hereafter: *the pragmatic thesis*).

When discussing this simplified version of contextualism, which comprises both semantic and pragmatic theses, we are going to confront it with some non-skeptical *invariantist* objections. Under the label "invariantism" we adopt and adapt the minimal non-skeptical anti-contextualist strategy, analogous to the semantic position articulated by P. Unger (1984/2002). According to it, the relevant standards for epistemic claims do not vary contextually as epistemic contextualists maintain. Consequently, the truth-value of epistemic claims and the epistemic position of the subject are not dependent on the features of the attributor context. Certain invariantist arguments point to the epistemic relevance of the *pragmatic* features of the *subject* context and of "what is at stake" in the subject's overall situation. These arguments come under the labels "subject-sensitive invariantism" by J. Hawthorne (2004) or "interest-relative invariantism" by J. Stanley (2005), having F. Dretske's (2000) externalist analysis of "the pragmatic dimension of knowledge" as their important predecessor in that sense. However, we do not take this pragmatic line and state our arguments in the purely truth-oriented or veritist terms: we consider knowledge, roughly, as a state functionally dependent on some truth-conducive factors like the processes producing sound inferences or other (conditionally) reliable processes of belief production. So, applying Stanley's (2005) and DeRose's (2009) labeling, our position would be then approximate to "intelectualist" or "classical invariantism".

Our invariantist objections to contextualism are grounded on our considerations of some of the contextualist arguments together with the fictitious cases and respective "intuitions" contextualists rely on. As it seems, our objections find also strong empirical evidence and a major methodological support in the field whose cognitive standards actually favor a truth-oriented and invariantist approach. The evidence is for a long time at hand: it derives from the well-known empirical findings in *cognitive psychology*, e.g. in the works of P.Wason, P. Johnson-Laird, D. Kahneman, P. Slovic, A. Tversky, or P. Cheng, J. Holland, K. Holyoak, R. Nisbett, P. Thagard, and many others. We are trying to show, at least in a programma-

tic way, that these results, which have already caused a tremendous impact in some philosophical areas, have significant bearings on the epistemic standards in general, especially concerning the status of claims attributing or denying *inferential knowledge*.

In that what follows we are trying to show why *changing descriptive* and *invariant normative* elements in inference practices and in their evaluations are important for inferential knowledge and respective knowledge claims. Fictitious cases of ordinary situations designed to support contextualism are shown not to provide unambiguous evidence for contextualism; moreover, slightly modified, they may be used equivalently as evidence for various kinds of epistemological *invariantism*. The upshot of our objections is that contextualism does not explain why reasoning subjects fail to have knowledge in the cases of fallacious reasoning and inferential ignorance described in psychological literature (like Wason's selection task and the conjunction fallacy). As a preliminary, we are going to point to certain *substantive* features of contextualism, beyond the semantic or pragmatic domain. These features present contextualism as the position more susceptible to a criticism from an invariantist point of view.

II. Contextualism and knowledge attributions

The central phenomenon which contextualism is intended to "save", i.e. to describe and explain adequately, is the existence of mutually exclusive, but nonetheless true and legitimate epistemic claims (where $p = S$ has hands, $BIV = S$ is a brain in a vat), nicknamed by DeRose (1995) as an "abominable conjunction", and construed as an argument for global skepticism:

(i) *S* knows that *p*.
(ii) If *S* knows that *p*, *S* knows that ¬*BIV*.
(iii) *S* does not know that ¬*BIV*

The common-sense realist accepts (i) and (ii), but not (iii), and the skeptic accepts (ii) and (iii), but not (i). However, the contextualist tries to reconcile these incompatible intuitions, maintaining that all these three claims are true, but in their own contexts. So, both "*S* knows that *p*" and "*S* does not know that *p*", which is derived from (ii) and (iii), are true, according to respective standards for knowledge attributions and to the use of the predicate "knows": the epistemic claim "*S* does not know that *p*" is true in the high-standards or skeptic context, while "*S* knows that *p*" is true in the low-standards or common-sense context. Due to the indexicality of

"knows", however, here we do not have a contradiction, since "knows" as used by the skeptic and "knows" as used by the common-sense realist express different contents, different epistemic properties. Contextualism is, therefore, presented as a solution to the problem of skepticism because it provides "the best explanation" of the genuine phenomenon of incompatible, but true epistemic claims.

Formulated at first as a response to global skepticism and brain-in-a-vat scenarios, contextualism now appears as a universal, although, in itself, slightly diversified strategy for dealing with various epistemological problems and paradoxes, like the *Lottery Paradox*, the *Epistemic Closure* and the *Gettier problem*.[1] Here, however, the focus is on the *generalized* contextualist strategy, which is not directed at solving the skeptical or any other specific problem as such, but to explaining the *context sensitivity of knowledge claims* in general. So, contextualists in general explain the differences in epistemic positions of the *same subjects* across *various attributor contexts* and *due* to the contextually shifting standards. Variations in the contexts of knowledge attributions and, consequently, in the epistemic position of the subject, follow the change of the *conversational standards* which govern the epistemic discourse.

The generalized contextualist strategy provides a finer gradation of various contexts and avoids extreme ranges of alternatives. Comparable to the case of incompatible, but true epistemic claims in the skeptical and non-skeptical contexts, contextualists, when distinguishing various contexts, similarly divide them into *higher-standards* contexts and *lower-standards* contexts. However, since their *Gedankenexperimente* imitate ordinary situations, in which ordinary speakers make epistemic claims attributing or denying knowledge to ordinary subjects, communicated in ordinary language expressions, naturally, *pragmatic* features come to the fore.

Like in the widely discussed *Bank Case* (K. DeRose 1992, and later) and *Airport Case* (S. Cohen, 1999), this gradation depends on the pragmatic features of the attributor's situation: the epistemic status of a believer is to be determined with respect to "what is at stake" in the conversational situation in which knowledge claims are made. So, the subject may know in a lower-standard context, when nothing vital is at stake, but fails to know in a higher-standard context, when ignorance could lead to some, say, financial disaster. Consider, for example, the *Airport Case* invented by S. Cohen (1999: 58):

[1] This diversification sometimes leads to disagreement: for example, S. Cohen (1998) and M. Heller (1999), as opposed to D. Lewis (1996), maintain that the Gettier problem is not to be solved by the contextualist explanatory means.

> Mary and John are at the L.A airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and respond, "Yes I know–it does stop in Chicago." It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, "How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute." Mary and John agree that Smith doesn't really *know* that the plane will stop in Chicago. They decide to check with the airline agent.

So, "what is at stake" for knowledge attributors Mary and John proves epistemically relevant when Smith's epistemic position is considered. Obviously, nothing changed in the subject context; however, new practical features in the attributor context raise the standards for attributing knowledge to the subject, so it seems to Mary and John that the claim "Smith knows that the flight stops in Chicago" is not true in this situation.

Such specific empirical or quasi-empirical findings based on the analysis of ordinary language utterances and epistemic "intuitions" are supposed to support contextualism. However, variations of the same cases, introducing additional elements, make the cases more ambiguous, so their evidential status becomes more susceptible to different and rivalling "intuitions". The variations of the cases invented for the contextualist purposes were used as evidence for their epistemological counterpart, *interest-relative invariantism*, as well. J. Stanley's (2005: 3–6) variations of the *Bank Case* include various combinations of respective stakes (*Low Stakes*, *High Stakes*, *Low Attributor-High Subject Stakes*, *Ignorant High Stakes*, *High Attributor – Low Subject Stakes*) and without a clear cut divides between subject and attributor situations. In the case where attributor stakes are low, and subject stakes are high, attributor's (Jill's) relaxed claim that the subject (Hannah) knows that the bank will be open on Saturday is "intuitively" *false* (J. Stanley 2005: 4):

> Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their pay checks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their pay checks by Saturday. Two weeks earlier, on a Saturday, Hannah went to the bank, where Jill saw her. Sarah points out to Hannah that banks do change their hours. Hannah utters, "That's a good point. I guess I don't really know that the bank will be open on Saturday". Coincidentally Jill is thinking of going to the bank on Saturday, just for fun, to see if she meets Hannah there. Nothing is at stake for Jill, and she knows nothing of Hannah's situation. Wondering whether Hannah will be there, Jill utters to a friend, "Well, Hannah was at the bank two weeks ago on a Saturday. So she knows the bank will be open on Saturday".

The attributor's (Jill's) claim attributing knowledge to the subject (Hannah) is false because it is not sensitive to the *subject context* and her practical interests. Were the situation in the subject context different and were the subject stakes lower, her epistemic claim would be, presumably, following Stanley's argumentation, true. Therefore, the slight change in the *Bank Case*, made by introducing a relaxed knowledge attributor and raising the stakes for the subject, turns the case against contextualism and in favor of interest-relative invariantism.

However, similarly to the *Airport Case* and the contextualist reading of it, we encounter the problem of the *relevance* of the knowledge attributor and of respective knowledge claims. So, we may pose a question, which is made from a trans-contextual perspective: "Why not, instead of Jill, to take Sarah to be the relevant attributor?" Or, in the *Airport Case*, "Why should we take Mary and John as relevant when Smith's belief is considered? Why don't we involve some trustworthy employee at the airport to check whether Smith knows?" In those revised situations we would have incompatible (Sarah's vs. Jill's) or possibly incompatible (an airport employee's vs. Mary and John's) epistemic claims. But, unlike the case of incompatible skeptical and common-sense epistemic claims, we would have obviously a clear way out: we, outside the particular attributor contexts, could pick out the relevant knowledge attributor – Sarah and an airport employee. The truth values of the pertinent knowledge claims would then be based on *truth-conducive* elements, like the reliability and the trustworthiness of the attributor or the quality of evidence, and not on stakes. For there is nothing at stake for an airport employee, who is, on the other side, obviously a more reliable attributor; and Sarah, although pragmatically dependent on the epistemic position of Hannah, is primarily relevant because she is in a more vantage evidential position than Jill. This is the point in which an *intellectualist contextualist* (as DeRose 2009 qualifies his position), and an intellectualist or *classical invariantist* could possibly agree in their "intuitions" on the cases.

But, at the same time the invariantist and the contextualist necessarily disagree on the cases: while the contextualist explains the epistemic position of the subject in terms of the sensitivity to *various* contexts and *various* contextual standards, the intellectualist invariantist halts the sequencing of contexts by making an attributor the *relevant* one and picking out the *relevant* standards and, therefore, uniquely determining the truth-conditions of epistemic claims and the meanings of the epistemic predicates.

The reason why the invariantist and contextualist interpretations of the cases are divergent may be of the *methodological* nature as well. The *pragmatic* considerations as well as the *linguistic* intuitions on ordinary

cases do their job in resolving the difficulties they are designed to resolve: primarily, to give a plausible explanation of the phenomenon of incompatible, but allegedly true epistemic claims. However, the invariantist would contest this contextualist pretense to solve serious epistemological problems, like that of skepticism, by means of the semantic considerations on the meanings of epistemic terms and on the truth-values of the epistemic claims. As E. Sosa (2000) claimed, the problem of skepticism is not just a problem of the correct use of the epistemic predicates and cannot be solved on the meta-linguistic level. After all, the skeptic challenges common-sense intuitions on knowledge on the *object* level, and therefore, needs an answer on the same level: it is not the question of the meaning of epistemic *predicates*, but of epistemic *properties* themselves.

However, as Cohen (1999) and DeRose (2009) rejoin, the semantic or meta-linguistic way is precisely the way out of the apparent paradox, which, according to their view, emerges on the "surface" of our ordinary epistemic claims, and should be resolved by semantic means. So, contextualism as an *epistemological theory* is not a theory of *knowledge*, but just a theory of knowledge *attributions*, a kind of semantic or meta-linguistic theory aimed at resolving paradoxes which are produced by the use of the linguistic devices expressing epistemic properties on the level of the epistemic object language. In his overview P. Rysiew (2011) conveniently positioned contextualism among *semantic* theories, meaning that it does not deal with *substantive* epistemological questions. Thus DeRose (2009: 21): the contextualism "is not a thesis about the structure of knowledge or justification", but is neutral to any such theory. However, although it does not give a structural specification of *what* knowledge is, contextualism obviously implies something substantive about knowledge (DeRose 2002: 168):

> Contextualists hold that the truth conditions of knowledge-ascribing and knowledge-denying sentences (…) fluctuate in certain ways according to the context in which they are uttered. What so varies is the epistemic standards that *S* must meet (or in the case of a denial of knowledge, fail to meet) for such a statement to be true. In some contexts, "*S* knows that *p*" requires that *S* have a true belief that *p* and also be in a *very* strong epistemic position with respect to *p*, while in other contexts, the same sentence may require for its truth, in addition to *S*'s having a true belief that *p*, only that *S* meet some lower epistemic standards.

According to the proclaimed semantic character of contextualism, the predicate "being in a very strong epistemic position with respect to *p*" should be read as epistemically neutral. For it allows various independent, non-contextualist and substantive construals. Some contextualists themselves in their non-contextualist moments put forward such construals,

like "being justified in believing that *p*" (by Cohen 1999) or "having a sensitive belief that *p*" (in the version by DeRose 1995, 2004), or "being in a position to exclude a relevant alternative to *p*" (by M. Heller 1999). But in the framework of contextualism as a theory of knowledge attributions, they restrain from these considerations. So, extending Kaplan's distinction to the meaning of the knowledge claims, the claim "*S* knows that *p*" *linguistically* and invariantly means that *S* is in a *good enough* epistemic position with respect to *p*, but that *how good* this epistemic position must be for *S* to count as knowing that *p* contextually varies, and this is its variable propositional *content* (DeRose 2009: 3). The question of content is, therefore, a *substantive* question: its solution does not lie within the scope of contextualism, because contextualists, by answering it, would commit the fallacy of "semantic descent" to the level of the object language, the language in which one talks *directly about knowledge*.

However, being indifferent as to *what epistemic property* is in question or *what degree* of its exemplification is sufficient for knowing, or *what kind of epistemic position* is denoted by this phrase, contextualism implies the fact that in any context *there is* an epistemic property exemplified in some degree and, therefore, that a subject *occupies* an epistemic position. It implies that, being the case that in different contexts different epistemic *propositions* are expressed, there is yet always some epistemic propositional content and epistemic property to be expressed. How otherwise could the strength of the subject's epistemic position be measured and graded as "strong" or "weak" if there would be nothing to measure and grade? DeRose (2009: 7) maintains that his concept of the strength of epistemic position is "entirely derivative from the concept of knowledge", meaning that it does not provide any informative or non-circular explanation of the knowledge concept itself:

> *To be* in a strong epistemic position with respect to some proposition one believes is for one's belief in that proposition to *have* to a high extent the property or properties the having enough of which is what's needed for a true belief to *constitute* a piece of knowledge. (our emphasis)

Obviously, *being in a strong epistemic position* is a function of the needed *epistemic property* or *properties exemplified* by the pertinent belief to a sufficiently high extent. We are inclined to read this DeRose's qualification as well as his statement that epistemic position derives from the concept of knowledge as *substantive*, at least in a modest or minimal sense of supposing that the epistemic properties, which are necessary for knowledge, are *in fact* exemplified (in a non-skeptical context) or not exemplified (in the skeptical context). Although the content of the knowledge concept and required epistemic properties are not specified by DeRose, this minimal

substantive sense comes to the fore in DeRose's quoted words (in italics), which emphasize factual belief's *having* epistemic properties as *constitutive* for knowledge.

If it is so, then the truth value of epistemic claims should be understood as dependent on the subject's epistemic position, which in turn supervenes on the epistemic properties actually possessed by the belief in question. And since "*S* knows that *p*" is true if and only if *p* is true and *S* is in a sufficiently strong epistemic position, the truth values of epistemic claims would derive from the "correspondence" between epistemic propositions and the pertinent epistemic "facts" (or, if some other notion of truth is operative by contextualists, the truth values of epistemic claims would derive from some other, perhaps, conventionally based relation between those claims and epistemic "facts" as exemplified epistemic properties or some other epistemic "truthmakers").

Of course, according to contextualism, the exemplification of epistemic properties is a matter of their attribution. Similarly, the truth value of epistemic claims is a matter of the attributed epistemic position according to the contextual standards. But then, reading these qualifications factually or realistically, as one obviously legitimately may do given the quoted sentences by DeRose, it would lead one to the idea that the contextually shifting standards and truth conditions of knowledge claims entail that the *truth value* of the epistemic claims changes. DeRose (2009) himself is cautious concerning the idea of the truth value change: it is rather that, depending on context and shifting standards for knowledge attributions, "*S* knows that *p*" and "*S* does not know that *p*" may be *both true*, but in different contexts, depending on different standards and in respect to different truth conditions. Yet, we think that it makes no essential difference: if *Ksp is* true in context C_1, and ¬Ksp *is* true in context C_2, then *Ksp* is *not* true in C_2 as well as ¬*Ksp* is *not* true in C_1. One way or another, the change of the truth value of the knowledge claims does occur, following the change of the truth conditions and standards for knowledge attributions. Proceeding that way, we are then already freely "descending to the object language" and "speaking of knowledge" (utilizing Rysiew's phrasing). We are forced to register not only the change of the truth values of epistemic claims, but to accept that the *fact* of the epistemic matter changes as well: if "*S* knows that *p*" is true in C_1, but not true in C_2, then *S knows* that *p* in C_1, but *does not know* that *p* in C_2.

So we would have, as perhaps an "unpleasant consequence", David Lewis' (1996) version of contextualism. According to it, not only truth conditions and truth values of epistemic claims fluctuate in dependence of the context, but the knowledge itself comes and goes as the context changes from a more relaxed to a stricter one: since "in the strict context

of epistemology we know nothing, yet in laxer contexts we know a lot" (1996/1999: 421), "it will be inevitable that epistemology must destroy knowledge. That is how knowledge is elusive. Examine it, and straightway it vanishes." (1996/1999: 435). Knowledge is, as Lewis then maintained, the function of relevant alternatives for *p* which are properly ignored by the knowledge attributor: for example, in more relaxed contexts, the knowledge attributor may properly ignore and eliminate the brain-in-a-vat alternative; but in the epistemological or skeptical context, the knowledge attributor may not ignore the brain-in-a-vat alternative since it is the relevant one. Consequently, in that context the truth conditions of knowledge attributions cannot be fulfilled and knowledge becomes impossible: it simply vanishes. Or, like in DeRose's caricature expression of it, "now you know it, now you don't".

Not all the contextualists would subscribe to this substantive and, at the same time, nihilist view (even Lewis (1996) himself in the last paragraph confessed that these "informal" statements should have been metalinguistically phrased). It is, however, a vivid illustration of the ambiguity of contextualism concerning the subject's epistemic position and its role in the truth value of knowledge attributions. DeRose (2009: 204–206) rejects this "now you know, now you don't" picture: a contextualist knowledge attributor may not say this, given her actual perspective and actual higher standards. She may not say that *S* does not know *now* (according to the actual higher standards) and did know *then* (before the standards got raised), because she can cite only *actual*, not *then* standards for epistemic attributions. But this maneuver appears problematic, since it leaves the obvious change in subject's epistemic position unexplained and does not save contextualism from the "now you know, now you don't" objection. If we follow the legitimate factual or realist reading of the relation between epistemic claims and exemplified epistemic properties, *something* obviously changed and that change should be observable to the same or to several knowledge attributors.

After all, both contexts (here: situations at different times) are connected by the same *character* of the predicate "knows", so the difference in *content* could be and should be recognized by the same attributor or be communicable between two attributors. This difference could then be traced back to some factual change of the epistemic position of the subject. In other words, the attributor A_1 could truthfully say to *S*: "Now, after the standards have been raised, you don't know!", implying that *S*, before the standards got raised, did know, since her epistemic position was dependent on *then*, more relaxed standards for knowledge ascription. Also, in a scenario with two attributors, A_1, the present attributor, could communicate this to A_2, the then attributor, because they share the same

linguistic meaning of "knows" and can agree on variations in the contents of knowledge attributions occurred after the standards got raised. And A_2 could say, "Yes, *S* knew that *p* according to then standards, but now, after the standards got raised, *S* doesn't know; I mean, I was not wrong in ascribing her knowledge *then*, since the standards then were not that high as *now*!"

In this particular case DeRose (2009: 185–225) appears as he himself "semantically descended" to applying a more substantive or objective-language term "*counts as knowing*", used by him as an equivalent for "*it can be truthfully said that the subject 'knows'*". So, according to DeRose,

"*S* counts as knowing that *p* in *C*"

is, therefore, a simplified translation for

"'*S* knows that *p*" is true in *C*'.

Obviously, assuming the purely semantic or meta-linguistic character of contextualism, this is not defendable: counting or considering *S as* knowing amounts to the direct attribution of knowledge to *S*. This is not just asserting the truthfulness of the knowledge claim "*S* knows that *p*". Or, to put it differently, the first one is ascription *de re*, the latter *de dicto*. In our reading the latter implies the first, but this reading does not hold if we confine the scope of contextualism to semantic considerations and are careful not to cause the level confusion. Yet, DeRose tried to vindicate this phrase as an "expository short-cut", along the lines Lewis was supposed to do so, according to DeRose's interpretation, when using his metaphor "elusive knowledge".

In the introduction to his *The Case for Contextualism* (2009: 19) DeRose described the content of the first volume as "an exercise in the philosophy of epistemological language", thus describing his present position more acutely. He, however, also clearly announced "a larger role" for "substantive epistemology" in the second volume (still to be published). In the same vein J. Schaffer (2006: 87) in his contextualist rejoinder to subjective-sensitive invariantist arguments put emphasize on the "irrelevance of the subject", meaning that the subject's stakes are not important for her epistemic position. He, however, clearly manifested the substantive talk: "The *knowledge relation* is sensitive to what is in question for the attributor, rather than what is at stake for the subject. There is no substitute for the contextualist semantics." (our emphasis)

Obviously, contextualism is not just a semantic or meta-linguistic theory of knowledge attributions, but a theory of knowledge as well. On DeRose and other contextualists then lies the burden of showing how this substantive move, "semantic descent", could be made in the framework

of meta-linguistically confined contextualist analysis, which dominantly develops in the direction of "semantic ascent".

Summing up, we can preliminary conclude that the contextualist program and methodology provoke the need for an alternative approach: the alleged, but not purely semantic character of contextualism and the ambiguous evidence of contextualist thought-experiments make the truth-oriented or intellectualist invariantist alternative a more plausible and more coherent view. Another point of invariantist dissatisfaction would be contextualists' seeking the evidence for their arguments in the "ordinary language basis". Indeed, why would one prefer ordinary language speakers as reliable or at least relevant knowledge attributors? Why think that these ordinary language epistemic claims are indicative of the epistemic positions of subjects and of knowledge at all? And why think that "knowledge", albeit indexical like "I" or "now", or gradable like comparative adjectives "tall" or "rich", is open to various, equivalently vindicated uses? The fact that people in their ordinary situations use "knows" differently does not justify our concessions to *all* their uses. In short, we think that the proper basis for deciding whether "knows" is indexical, or whether knowledge attributions are context sensitive, and if they are context sensitive, how they are sensitive, lies in a more stringent framework. This framework is provided in the substantive epistemological and psychological research on truth-conducive elements of knowledge, as our considerations in the next section suggest.

III. Inferential knowledge attributions and the psychology of reasoning

Our opting for the truth-oriented invariantism is primarily due to the "intuition" that the shift of the truth conditions of epistemic claims, their truth values and, finally, the change of the epistemic position of the subject, depends on and is sensitive to the cognitive equipment, environment or history of the *subject* in question. Accordingly, the subject's epistemic position is not a matter of pragmatic or interest relative features of the situation in which some knowledge attributor is situated (like Mary and John in the *Airport Case*, or Sarah in the revised *Bank Case*), but primarily of the *truth-conducive* or *truth-preserving* processes or procedures which caused or grounded her beliefs. Yet, what is *epistemically* important may also be *pragmatically* important in a *subject's* life situation. However, this pragmatic importance depends on the factors which conduce or preserve the *truth* in the process of acquiring or justifying her belief, and not vice versa. Many cases of propositional knowledge discussed from a subject-sensitive invariantist perspective, endorsed notably by F. Dretske

(2000), or J. Stanley (2005), however, illustrate this "pragmatic dimension" of knowing. Nevertheless, what proves more fundamental are the aforementioned truth-generating or truth-preserving processes. Therefore, even in the cases where pragmatic features of the subject's environment essentially influence the epistemic position of the subject, this fact is due primarily to her ability to produce relevant true beliefs.

Instead of the idea of interest-relativity, we adopt the idea, formulated and advanced by A. I. Goldman (1986, 1992, 2012) in various settings, from the individual to social epistemological considerations, that the epistemic position of a subject depends primarily (or exclusively) on her chances to achieve or produce true beliefs. Here, *veritism*, broadly conceived, is favored to *pragmatism*, since the truth is epistemically prior to interests. The reason is simple: for the realization of her own interests one necessarily relies on having true or at least empirically adequate beliefs, which could amount to the knowledge of facts. Otherwise, when the opposite scenario would hold, in which the epistemic position of a subject would be a matter of "what is at stake", not only the truth and knowledge would be in danger, but the fulfillment of her interests, too. In such a scenario either (i) the practical importance of a situation could (wrongly) deprave the subject of her knowledge, which she, in fact, has, or, (ii) the relative practical unimportance of someone's belief could turn her actual ignorance into knowledge. Both alternatives lead to disastrous epistemic and pragmatic consequences as well.

The truth, truth-conducing and truth-preserving procedures and processes are usually associated with *reasoning* and *inferences*. In epistemological terms, reasoning and inference yield *inferential justification*, which in turn can give *inferential knowledge*. An item of knowledge is inferential because the justification in question or the pertinent belief-producing process is inferential: *S* believes truly that *q* on the basis of *S*'s belief that *p*, where *p* is a set of premises explicitly entertained by *S* as her reasons for her belief that *q*. Inferential knowledge is, however, not necessarily derivative in this sense, since its justification element does not have to be internally accessible to *S*, or *S* is not necessarily aware of the quality of her premises and the validity of their inferential relations. The fact that a reasoning process is (conditionally) reliable does not have to be internally accessible to *S* in order to have an inferentially *justified* belief and inferential knowledge. This is the point we owe primarily to A. I. Goldman (1979, 1986) as a paradigm example of the *reliabilist* analysis of knowledge.

This approach is, however, comparable with the aforementioned *intellectualist* brand of contextualism. According to the intellectualist contextualist, as presented by DeRose (2009: 188–190), whether *S* may be credited with knowledge depends on how the attributor assesses the truth-

related factors, like the reliability of belief-forming processes at *S*'s disposal, or the sensitivity of her beliefs. This parallelism is also manifest in the contextualist main cases. For example, when assessing their own epistemic status as the state of knowledge or ignorance, Mary and John in the *Airport Case* do not *base* their epistemic judgment on their stakes. We are inclined to say that the pragmatic elements in their situation only *trigger* more fundamental considerations, those about the *reliability* of their *evidential* sources (Mary and John ask: "How reliable is that itinerary?") and their *procedures* ("They decide to check with the airline agent"). Similarly, in various versions of the *Bank Case* attributors' higher stakes *caused* the revisions of their initial knowledge claims. Yet, their epistemic judgments in the low as well as in the high stakes contexts were *based* on attributor's assessments of how *reliable* the subjects were in picking relevant information and in forming true beliefs, or how sensitive their beliefs were, about the bank working hours ("I'd better go in and make sure", as has been said by the husband, the first-person knowledge attributor; cf. DeRose 2009: 2).

This point is also expressed, ambiguously, but explicitly by DeRose (2009: 188):

> For the contextualist, exactly which proposition gets expressed by a knowledge-ascribing sentence will often be affected by 'practical' factors, but the particular proposition that does get expressed will not itself be at all about those factors: *Whether that proposition is true is determined only by the subject's attitude and the truth-relevant factors of the subject's situation*. (our emphasis)

We understand this point, however, as provoking the need for the *subject-sensitive* assessment of the truth relevant factors. According to DeRose, these factors are operative in "the subject's situation". After all, whether knowledge claims are *true* or *not* depends on them as well as on "the subject's attitude".

The following point made by DeRose's (2009: 188) discloses the ambiguity of his position even more:

> The contextualist does not hold that whether a subject *knows* or not can depend on non-truth-relevant factors; he holds that whether a speaker *can truthfully describe the subject as 'knowing'*– whether, in our sense, the subject 'counts as knowing' in the speaker's context– can depend on such factors. Whether the speaker can truthfully describe the subject as 'knowing' can depend on such factors, according to the contextualist, because such factors can affect the precise content of the speaker's claim, *not* because they can affect whether the subject is such as to make true the proposition that the speaker is asserting about her.

Here, obviously, DeRose maintains that the pragmatic elements *can* affect the content (i.e. the proposition) expressed by a knowledge ascribing sentence, and consequently, the *truth conditions* as well as *truth values* of the respective knowledge ascribing sentence. But we would rejoin: a knowledge ascribing sentence is true *because* the *proposition* expressed by it is true, and that proposition is true, as DeRose himself says, because of the *truth-relevant factors* operative in the *subject's* situation. Consequently, whether the knowledge claims are true or not is entirely dependent on the subject's situation, not on variable contexts of knowledge attribution.

Let us assume then that what contextualists, and, particularly intellectualist contextualists say about knowledge attributions is pertinent to the ascriptions of *inferential* knowledge as well.[2] We focus on the *general* problem of knowledge attributions, in the framework of which the problem of the inferential knowledge attributions is supposed to be just a *special* instance of it. When the emphasis is on the inferential knowledge attributions, the way the pertinent inference confers the warrant/justification to the conclusion believed is *crucial* for the epistemic assessment of the belief in the conclusion. Therefore, the attribution of inferential knowledge and any theory of it must pay attention to inferential or reasoning standards, by which inferential justification and knowledge are assessed. The inferential knowledge attributions are correlated with the strength of *S*'s epistemic position with respect to *q*, which is inferred from *p*. As it has been shown concerning the contextualist cases (*Airport Case*, *Bank Case*), *S*'s epistemic position with respect to *p* depends on (the assessments of) the *reliability* of the pertinent belief forming process or of the *sensitivity* of belief that *p*. Analogously, the (assessed) *validity* or the *cogency* of inferential processes from *p* to *q*, as pertinent belief-forming processes, should prove decisive for *S*'s epistemic position with respect to *q*.

From our perspective, therefore, *inferential standards* prove constitutive for, and actually function as, *epistemic standards*: whether the epistemic subject may be credited with inferential knowledge depends on the quality (validity, cogency) of her reasoning assessed by inferential standards. According to contextualism, for sure, the truth conditions of

[2] When doing so, we neglect the notorious epistemological problems related to the role of deductive or probabilistic inference in structuring knowledge, like the *Epistemic Closure*, the *Gettier problem* or the *Lottery Paradox* to mention the central ones. Significant attention has been paid to these problems among contextualists, for example, by S. Cohen (1998) D. Lewis (1996) and M. Heller (1999), as well as among subject-sensitive invariantists, by F. Dretske (1970, 1971, 2005) or J. Hawthorne (2004). However, these problems primarily arise when one poses substantive questions on *the structure of inferential justification and knowledge*. And it is still not clear whether the contextualism vs. invariantism controversy is particularly relevant for their solutions.

the sentences of the type "*S* knows (inferentially) that *q*" vary with the contexts in which attribution of inferential knowledge is made and according to its standards, and not with the contexts of the person who draws the inferences and comes to know that *q* inferentially. Yet, the norms of respective "conversational contexts", which govern the use of inferential knowledge claims, simply *cannot* avoid the reference to the standards of logic and probability calculus, dominantly deemed as *invariant* in a relevant sense. As it seems, the invariability of these standards should be accepted by contextualists, too. And if it is so, then *their invariant epistemic roles* pose significant difficulties for contextualists.

Accordingly, the massive empirical evidence coming from *cognitive psychology's* subfield psychology of reasoning, as it seems, is not favorable to contextualism. This evidence contains research data concerning the reasoning performances of individuals in various types of situations. The works of P. Wason (1966), P. Wason and P. Johnson-Laird (1971), D. Kahneman, P. Slovic and A. Tversky (1982), A. Tversky and D. Kahneman (1984/2002), J. Holland, K. Holyoak, R. Nisbett and P. Thagard (1986) and many others collected and systematized experimental data in the span of several decades. These empirical results have already been utilized extensively in philosophical discussions on human rationality and reasoning capacities as well as in epistemological and meta-epistemological considerations on the nature of inferential knowledge and the proper philosophical methodology: notable examples include A. I. Goldman (1986, 2012), H. Kornblith (1993), S. Stich (1990), to mention just a few. We believe that these empirical results and their evaluations are highly relevant for discussing the status of knowledge claims and the epistemic position of the subject as well. The way the reasoning performances of "real people" are evaluated and the way their outcomes – the pieces of *inferential knowledge* or *inferential ignorance* – are attributed to the reasoning subjects may prove illuminative when one comes to the question of the truth conditions and the truth values of the knowledge claims of the form "*S* knows inferentially that *q*".

Cognitive psychologists demonstrated that people in various contexts mostly do not reason according to logical or probabilistic rules, but according to other standards, the reliance on which leads to systematic logical or probabilistic errors. D. Kahneman, A. Tversky and others identified certain *heuristics*, some other authors from the field suggested that people in fact follow different reasoning rules they called *pragmatic reasoning schemas* (J. Holland, K. Holyoak, R. E. Nisbett and P. Thagard 1986) or *domain specific rules* (L. Cosmides 1989) or *mental models* (P. Johnson-Laird and R. Byrne 1991). However, independent of the strategies *actually* applied in everyday or expert reasoning practice, the general point which

can be made and actually *has been made* on the basis of the psychological research in the domain of the psychology of reasoning is that reasoning subjects, in a wide range from lay to expert contexts, make inferential and probabilistic fallacies with respect to contextually *invariant standards*, or, at least, the standards which are treated as invariant from a particular *expert* point of view. They are not always the standards which subjects in fact *follow*, but the standards by means of which their cognitive performances are relevantly *assessed*. How otherwise could their performances be assessed as *"fallacious*" or as leading to "*systematic errors*", as psychologists often emphasize?

These psychological findings concerning human reasoning performance undermine contextualism, at least, indirectly. Obviously, individual reasoners frequently do make fallacies according to invariant inferential (deductive or probabilistic) standards, and *therefore* fail to gain inferential knowledge, but not because knowledge attributors in those cases change their standards. As those findings indicate, the standards did not shift, but certain variations in subject contexts, like domain specificity of reasoning cases or the influence of pragmatic features, essentially influenced their inferential performances.

The assessment of inferential performance, and its result, inferentially produced beliefs, has to rely on invariant standards by means of which it is established what the correct inference is. For an invariant normative level determines or confines the phenomenon which is the object of investigation at the descriptive level. Without a general idea about what correct inference is we would hardly define what in general (any) inference is. Also, studying the way reasoners *actually* reason is not possible if one does not have an answer how reasoners *should* reason.

Our point here is that the truth conditions of the claims attributing or denying *inferential knowledge* similarly depend on invariant standards, which involve inferential or probabilistic rules as their essential part as well. So, when trying to explain why the sentence "*S* knows inferentially that *q*" may vary and be true in the context C_1, but false in the context C_2, we treat the contexts C_1 and C_2 as the contexts in which *S* comes to believe that *q* and makes some inference from *p* to *q* according to some inferential rules, not the contexts in which knowledge claims are made. The variations in *S*'s overall epistemic position are explained as the function of the facts in *her* context, and not in the context in which her position is assessed. Or, to put it in another way, the fact that *S* knows that *q* in some overall situation or context C_1, but fails to know that *q* in a different overall situation or context C_2, is not a consequence of the changing standards for knowledge attribution– since they do not change or their change is not relevant – but of the changing facts about her cognitive performance or epistemic envi-

ronment. Consequently, it is proved that the variable truth conditions and truth values of knowledge claims are sensitive to the changing facts in the subject context. We find compelling reasons against contextualism in that empirical realm. In that way we do not only question their *semantic* and *pragmatic* theses, but at the same time contrast their intuitional and linguist methodology oriented on collecting "data" from contextualist thought-experiments, with the pieces of genuine empirical evidence.

In order to "test" contextualism in the realm of inferential knowledge, let us first consider two typical and widely discussed cases of formally incorrect reasoning, described in psychological literature. The first is the case of reasoning in the framework of famous Wason's *selection task*, presented in numerous versions by P. Wason, P. N. Johnson-Laird, R. E. Nisbett and others to show how logically untrained people are biased to make mistakes in conditional reasoning. The second is the paradigmatic case of erroneous probabilistic reasoning, the *conjunction fallacy*, reported by D. Kahneman, A. Tversky and others.

1. *Selection task*

In the original version of the selection task, four cards are presented to subjects.

E K 4 7

Every card has on its uppermost face a single symbol. Subjects knew that on one side of each card is a letter, and on the other a number, according to a rule. Then the conditional was presented:

> If a card has a vowel on one side, than it has an even number on the other side.

Subjects had to select which cards are to be turned over in order to establish whether the conditional is true or false. The correct selection was the E card (which must reveal an even number on the back if the rule is true) and the 7 card (which must not reveal a vowel on the back). This combination, a vowel on the one side and an odd number on the back, renders the case which corresponds to the case when the antecedent is true and the consequence is false; and this is the only case in which this conditional is false, according to the truth table for conditionals.

When we express the reasoning they should have performed or the reasoning a competent subject or the inferential knowledge attributor would perform, in order to test the rule, we get, roughly, the following conditional reasoning schema:

If a card has a vowel on the one side, then it has an even number on the other side.

$(p \rightarrow q)$

It is not the case that, if a card has a vowel on the side, then it has an even number on the other side if and only if the card has an odd number on the one side and a vowel on the other side.

$\neg (p \rightarrow q) \leftrightarrow (p \wedge \neg q)$

When negative evidence occurs (when it happens that the card has vowel on the one side and the odd number on the other), it conclusively disconfirms the rule, so we have an instance of modus tollens:

$(((p \rightarrow q) \rightarrow \neg(p \wedge \neg q)) \wedge (p \wedge \neg q)) \rightarrow \neg(p \rightarrow q)$.

The reasoning subject must predict the behavior of the cards according to the truth-table for material conditional. So, this correlation between the different values of the antecedent and consequent on the one hand and the factual relations of cards on the other make this task a problem in the real world: the logical relation functions as a clue for discovery of a real world (albeit local) regularity. Accordingly, the insight in those formal relations is a necessary justificational element in the knowledge of these factual relations. Initially posed as a formal test of whether people do reason correctly, the selection task is also a test of the ability of people to gain inferential knowledge of the fact or content covered by this rule.

An insignificantly small number of subjects in the original and subsequent experiments selected the right cards, not more than 20%. The majority of subjects turned the E card or the E and 4 cards, which demonstrates the pervasiveness of the *confirmation bias*, i.e. the tendency to seek primarily positive evidence and verification instead of falsification. Similar results followed in other such, relatively formalized and non-contextualized versions of the task, when the cards had (i) other combinations of vowels and numbers, (ii) circles on one side and blank surfaces on the other sides, etc.

The situation changed when the contexts of examination and the analogical selection tasks became more familiar and less formal, closer to subjects' everyday situations, e.g. when the material conditional has been shaped as a rule for socially proper conduct or social contract. So, if we, instead of semantically poor tasks with numerals and letters, or geometrical figures, formulate a task which refers to some social rules, like

If a person is drinking beer, then the person must be over 18.

or to a kind of social obligation, like

If you take my car, you should fill up the tank,

then, what is called "the content effect" (P. Johnson-Laird and R. Byrne 1991: 73) or "the framing effect" influences conditional reasoning: it became *more successful* with respect to the conclusions drawn. In such contexts subjects reached the right conclusions and produced true beliefs about what cards are to be drawn in order to test the rule to a much greater extent. For example, in the permission schema the task was to test the rule that, if a person is drinking beer then the person must be over 18, the falsifying case was the combination "the person is drinking beer" and "the person is not over 18" cards.

An influential explanation of the content effect has been proposed by P. Cheng and K. Holyoak (1985/2008: 829) in terms of "pragmatic reasoning schemas":

> We propose that people often reason using neither syntactic, context-free rules of inference, nor memory of specific experiences. Rather, they reason *using* abstract knowledge structures induced from ordinary life experiences, such as "permissions," "obligations," and "causations." Such knowledge structures are termed *pragmatic reasoning schemas.* A pragmatic reasoning schema consists of a set of generalized, *context sensitive* rules which, unlike purely syntactic rules, are defined in terms of classes of goals (such as taking desirable actions or making predictions about possible future events) and relationships to these goals (such as cause and effect or precondition and allowable action).
>
> Our theoretical framework assumes that the role of prior experience in facilitation is in the induction and evocation of certain types of schemas. Not all schemas are facilitating, as becomes clear below. Some schemas lead to responses that correspond more closely than others with those that follow from the material conditional in formal logic. Performance as *evaluated by the standard of formal logic* depends on what type of schema is evoked, or whether any schema is evoked at all.

We find this explanation by Cheng and Holyoak strikingly relevant for our previous considerations on the context sensitivity of knowledge claims. It is clear that Cheng and Holyoak see reasoning performances as pragmatically and contextually conditioned by the reasoner's *use* of pragmatic reasoning schemas as "*context sensitive* rules" defined in respect to certain pragmatic "goals", i.e. to "what is at stake" in subject context. By using proper reasoning schemas subjects reached better inferential results, i.e., they reproduced more frequently a reasoning schema which corresponds to *material conditional* in *formal logic*. On the other hand, the inference failure, and therefore inferential ignorance, in the original, non-contextualized version of the selection task is due to the fact that it did not evoke a reasoning schema corresponding to *material conditional*.

Obviously, reasoning performances are "*evaluated by the standard of formal logic*", so the pertinent inferential knowledge or ignorance can be claimed only on the background of this authoritative *invariant standard*. We could say, in a slightly rigorist way, that "*S* knows inferentially that *q*" is true or false according to the context insensitive or invariant standard. It means that people gain respective inferential knowledge insofar their reasoning structurally corresponds to relevant logical rules, and do not gain it insofar they do not reason according to such rules. Pragmatic reasoning schemas enable better reasoning results, but only because reasoning performance, statistically, tend to yield formally correct inferences.

As Holland, Holyoak, Nisbett and Thagard (1986) report, in such an experiment, 66% of subjects which were confronted with a "permission problem" like the previous "drinking age" rule or a similar, contextualized version of the task got the right answer versus only 19% of subjects which ought to solve an arbitrary and formal problem similar to the original selection task. This experimental fact was reproduced in similar ratios in many other experiments.

All these research results confirm our initial assumption that differences in cognitive performance and, therefore, in the *epistemic position* of the subject are not due to the variations in the relevant *attributor context* – since there are no such variations – but only to the variations in the *subject context*. These variations include different reasoning schemas used by reasoners and different pragmatic goals: goals ("what is at stake") determine the *choice* of the scheme, but the scheme is epistemically successful if it follows a standardized *norm*. Only in this way subjects, proceeding from true premises, $p_1, \ldots p_n$, non-accidentally or necessarily reach the truth and achieve the pertinent inferential knowledge of the consequence *q*.

Cognitive scientists identified, in fact, various cognitive instruments which have a facilitating effect on reasoning in the selection tasks, ranked according to their generality: from *deontic logic* as the most general one, which is proposed by D. Cummins (1996), and *pragmatic reasoning schemas* in terms of permissions or obligations, considered by Cheng and Holyoak (1985), over the domain of *social exchange* and the *social contract* algorithm by L. Cosmides (1989) to the "checking for cheaters" algorithm as the most concrete one by G. Gigerenzer and K. Hug (1992). Different levels, of course, do not have to exclude each other, so in one reasoning process the simultaneous activation of several levels is conceivable: deontic logical rules, pragmatic schemas, the kinds of content and specific experience. Unlike the theses explaining the content effect in Wason's selection task by the kind of content on which subjects make their inferences, D. Sperber, F. Cara and V. Girotto (1995) in the framework of their relevance theory explain the same phenomenon without evoking any

specific kind of content. The reasoning outcome depends on the reasoning context. However, the result is the activation of the general mechanism of the linguistic understanding. Information is assessed as relevant when it causes a cognitive effect, i.e. leads to forming a new belief or to rejecting the current one. A cognitive effect includes processing, which amounts to a cognitive effort. When interpreting one's assertions, a subject tries to *maximize the relevance*, and it means that she accepts that interpretation which, in particular circumstances, by means of economic cognitive effort, yields a sufficient cognitive effect. In the contextualized versions of the selection task (e.g. alcohol/age of majority) most of the examinees give a correct answer because the context of such tasks, unlike the original version, diminishes the cognitive effort of processing and augments the cognitive effect of the correct answer.

In spite of significant mutual disagreements, all these explanations refer to the subject context, or more general, to the structure of the environment. An environment contains elements or information which reasoning subjects judge as relevant for solving a particular problem, so they are included in the very process of reasoning and thus influence its outcome. This is obvious evidence for the thesis that the *difference in performance* of reasoning subjects and in their overall cognitive achievement is a matter of the *shifts in their cognitive surroundings* and the *quality of the processes* of their problem solving. Even if there is no correlation between the reasoning success and the intuition people share about logical rules, additional information and a richer evidential basis obviously contribute to the reasoning success of examinees, or at least, to the right guessing in subjects' everyday situations.

Such an improvement in reasoning outcomes of the subjects is not necessarily an indicator of their *epistemic improvement*, since the experiments show that the subjects sometimes guessed the right solutions through unreliable or improper procedures. We are inclined to say that, although they achieved better results after the tasks were rephrased, they were not in a principally better epistemic position concerning the task and its solutions. If we take a strong rigorist and intellectualist invariantist attitude, we would say that, basically, they remained *ignorant* about the solution in question. For the acquaintance with the truth table for conditional and its right application as well as a competent application of elementary reasoning rules, such as modus ponens and modus tollens, are *necessary* conditions for the knowledge of the conclusion. Any right guesswork simply would not count as knowledge.

But to deny them *any* epistemic improvement would confront us with an obvious change in their performance. Since they give *locally* correct answers, i.e. right solutions to *informal* problems, and not to purely formal

logical problems, they have a kind of local "knowledge" according to the conditions in the subject's environment. Even when subjects reason in a formally correct way, logic does not determine which of the infinitely many valid inferences they will perform in a particular case. This is determined by the pragmatic features of the situation. Logic discerns valid from invalid inferences, but particular circumstances determine *which* valid inferences, among infinitely many possible valid inferences, are useful or relevant in these circumstances. So the epistemic improvement, as well as the epistemic degradation, is in that sense a matter of pragmatic or situational factors operating on the fixed normative background.

Contextualists would say that whether the subjects know, or in metalinguistic phrasing, whether "S knows that *q*" is true, in such situations depends on the contextually variable attributor standard applied when the results are assessed. The measuring standard applied here is normatively fixed: it is connected to indicative conditionals and classical bivalent propositional logic. It is not any of the standards which the examinees perhaps implicitly applied, e.g. heuristic principles or some other contextually or domain specific reasoning schemas. For only by means of such an *invariant* standard could reasoning fallacies or improvements get detected. In fact, empirical research was designed in a way to check whether logical laws, which constitute one of the standards of inferential knowledge, hold in practice as reasoning rules. The negative results were interpreted as a deviation from *them*. Cognitive psychologists tried to establish and to explain the nature of the relation between inferential practice and logical rules. Cognitive shortcuts used in reasoning, environmental conditions of their successfulness, kinds of content which have a facilitating or obstructing effect on reasoning – all these phenomena have been detected and assessed only by referring to the invariant standard of the logical norm. This is strongly indicated by the *convergence* of results toward the same value in various contexts as well as by their parallel "*improvement*" when the tasks got conveniently contextualized.

In order to illustrate it, we may apply the contextualist strategy and the idea that "knows" in different contexts of examination means different things. Suppose that there are two groups of examinees, one which is formed by the graduates in class A, in the philosophical department where a kind of mystical existentialism is favored, and the other, formed by the graduates which are trained in logic and are familiar with formal epistemological problems (like brain-in-a-vat); they are in class B. The task to be solved is the original selection task. Students are also required to explain their conclusions, but the criteria for this explanation are higher in class B, than in class A: the first should refer to the truth table for material implication, the second should give any explanation whatsoever.

However, a student *S* from class B happens to solve the task among the members of class A. After checking the results, the teachers notice the mistake and her test is now together with the tests of her colleagues in class B. At first, they assessed her result as correct in the given margin for class A: she has chosen the right combination of cards, but explained her choice by referring to the application of some inappropriate formal rule and to the truth table for, say, material biconditional, and not by referring to the complete truth table for material conditional and to modus ponens and modus tollens. When assessed among the tests from class B, *S*'s result is not so satisfactory any more: her explanation is not good enough for those teachers. Did *S* know in the first case, when her test was checked in the context in which standards for the class A were operative, but did not know when they were checked according to the criteria for the class B? Or, in other words, is the epistemic claim "*S* knows that *q*" [i.e. the right combination of the cards] when uttered in the first context true, and, when uttered in the second context, false?

If this is the case, shall we say, then, that "knows" in both contexts denotes different epistemic properties, so that in the first case, in order to reach the right answers and acquire knowledge, one must meet lower standards, and that in the second one must satisfy higher standards? Hardly! In both contexts, even if the teachers of both groups applied different epistemic norms/criteria, we would expect the same performance in order to ascribe to *S* the reasoning success without the danger of systematic error and in order to attribute to *S* knowledge about the conclusions drawn: the *same cards* should have been picked up according to a single logical standard. Pure guesswork would not have worked: only the knowledge of the truth table and valid reasoning schemas would have given the knowledge of the conclusion (given enough relevant information), making the claim "*S* knows that *q*" is true.

Consider the previous case, but in the version in which the epistemic position of the subject is correlated with the pragmatic features of the attributor's situation. For example, imagine that *S* has a true belief that *q*, about which cards are to be turned over, yet reached by a wrong inference rule, and that *A*, some attributor, is to lose ten cents for her wrong answer: Is the attributor's claim "*S* knows that *q*" then true? Also, imagine some other context, in which the attributor could lose a million euros for *S*'s wrong answer, and *S* has an inferentially justified true belief that *q*, immune to Gettier counter examples: Is the claim "*S* knows that *q*" false because of the stakes in the attributor's context? It is hard to see why the respective knowledge claim in the first case would be true and in the second false. In both cases the measuring standard and the normative expectations concerning the result stay the same, even when there are some

differences in respect to what is at stake for the attributor. In both cases *S* would be right or wrong and she would know or would not know which cards are to be turned over irrespectively of the stakes. So, when all factors are doing their truth-conducive job properly, "*S* knows that *q*" is true, and when these factors are not truth-conducive, "*S* knows that *q*" is false, no matter what is at stake for the attributor.

2. *The conjunction fallacy*

One could find previous considerations on the status of knowledge claims in the framework of deductive reasoning and respective inferential knowledge perhaps not so persuasive or interesting, since the cases of inferential knowledge or ignorance based on deductive inference are not just cases of *factual* knowledge or ignorance, but primary of logical knowledge, or knowledge of logical connectives and relations. We find this complaint unjustified, since in all versions of the selection task the knowledge of the logical rules or relations was only instrumental in discovering *factual* regularities. Nevertheless, this complaint would perhaps not be issued to the following considerations on the probabilistic reasoning and inductive inferences, since they produce beliefs whose factual content is undisputable, although similarly dependent on other type of *formal* rules, i.e. the rules of the probability calculus.

In many everyday situations probabilistic inferences we draw are, however, not a result of the application of the standard norm, i.e. the probability calculus. In everyday circumstances one often lacks the time or information, or numeric values required for the application of the norm. In the real process of probabilistic reasoning, therefore, instead of very complex calculations prescribed by the norm, people use cognitive shortcuts, so-called *heuristics*. By means of them they reduce the extension of the problem, thus making the probability judgment simpler. For example, the *availability* or associative distance is a heuristic in which subjects, when judging the probability, rely on several strategies of "providing" the information assessed by them as relevant for solving the problem: when reasoning about new features in the situations subjects usually retrieve familiar features from the memory. *Anchoring* and *adjustment* present the dependence of the assessment of some quantity on a previously available quantity, and the heuristic of *representativeness* applies to various inferential strategies in which the probability of an outcome is judged on the basis of the degree to which it resembles the essential properties of the evidence. For *resemblance* is, unlike the probability, an easily available and easily assessable value.

In ordinary circumstances relying on heuristics mostly leads to sufficiently good results: it is easier to recollect those events, which in ordinary circumstances are indeed more frequent; it is easier to imagine or to construe events which are more likely to happen, and associative links strengthen when two types of events frequently happen together. In all these cases inference to the frequency, likeliness and correlation by means of the heuristic of availability may be successful. The effect of anchoring is correlated with cognitive economy, and the heuristic of representativeness is defendable by the fact that in many circumstances more probable events are more representative than those less probable. However, since heuristics are to a great extent based on accidental and subjective factors in the contexts of reasoning, they may result in systematic fallacies, so-called *biases*. For example, the heuristic of representativeness results in expectations that any two samples picked out from the same population resemble each other and the population more than it is predicted by the sample theory, namely, it is expected that any segment reflects the real proportion (so-called local representativeness) and that a random process is auto-corrective.

In the preface to a retrospectively composed collection of essays devoted to the discovery of "heuristics and biases", which according to the authors actually govern people's reasoning, the editors T. Gilovich and D. Griffin emphasized the contextual and situational relativity of people's reasoning performance. At the same time, they restated the obviously entrenched and widespread *invariantism* concerning the standards to which this performance has been assessed:

> The core idea of the heuristic and biases program is that judgment under uncertainty is often based on a limited number of simplifying heuristics rather than more formal and extensive algorithmic processing. These heuristics typically yield *accurate judgments*, but can give rise to *systematic error*. (T. Gilovich and D. Griffin 2002: xv) (our emphasis)

A. Tversky and D. Kahneman (1984/2002) in their study on extensional versus intuitive reasoning again forcefully express the same idea:

> The presence of bias and incoherence does not diminish the *normative force* of these principles, but it reduces their usefulness as *descriptions* of behavior and hinders their prescriptive applications (A. Tversky and D. Kahneman 2002: 48) (our emphasis)

The contingent and pragmatically justified "accuracy" of these heuristics obviously cannot obscure the fact that they, when the context and content of reasoning are changed, and it may always happen, lead to a "systematic error". And, although the cognitive standard is not *descriptively* and

prescriptively useful, it obviously stays as an invariant *evaluative* standard for the attribution of inferential knowledge in this domain. Obviously again, the differences in cognitive successfulness are due to the changes in cognitive environment or in the subject context, while the attributor's standard remains the same.

A paradigm case of the "judgment under uncertainty" is the so-called *conjunction fallacy*, which illustrates how the content and context influence reasoning results. Consider the famous case:

Linda is 31 years old, single, outspoken, very bright, majored in philosophy, interested in social and moral issues, active in anti-nuclear movement. So, we have the following descriptions of her:

1. L. is a teacher in elementary school.
2. L. works in a bookstore and takes Yoga classes.
3. L. is active in the feminist movement. (*f*)
4. L. is a psychiatric social worker.
5. L. is a member of the League of Women Voters.
6. L. is a bank teller. (*t*)
7. L. is an insurance salesperson.
8. L. is a bank teller and is active in the feminist movement. (*f*&*t*)

The most frequent result (85% of the examinees) was the answer that 8. (*f*&*t)* is more probable than 6. (*t*). This answer clearly violates the rule for probabilistic inferences that conjunction cannot be more probable than its conjuncts. If *I* = background information about Linda, *f* = Linda is active in the feminist movement, and *t* = Linda is a bank teller, then their erroneous inference in question goes in the following way:

Given *I*, $P(f) > P(t)$
Therefore, $P(f\&t) > P(t)$.

The conjunction fallacy is caused by the fact that the logic of representativeness differs from the logic of the probability calculus: additional elements in a description may increase the representativeness of an event or its resemblance to a stereotype, however, without increasing its probability. In short, although it cannot be more probable, conjunction may be more representative than any of its conjuncts. So, in such circumstances, the assessment of probability based on the heuristics of representativeness leads directly to the conjunction fallacy.

How to interpret this empirical fact of the pervasive erroneous reasoning and inferential ignorance in the light of contextualist strategy of explaining the truth value of knowledge claims in dependence of the attri-

butor context? Taking semantic and pragmatic theses as contextualist starting points, two procedures are available: (1) to stratify the meaning of the pertinent "knows" concerning the proposition "P(*f*&*t*) > P(*t*)" in relation to various conversational contexts, from the lay to the expert ones, and (2) to vary the epistemic standards which govern knowledge attributions in those various contexts according to "what is at stake" for the attributors.

However, as in the previous case of conditional reasoning, there is not much space for contextualist maneuvers. Since contextualism is focused on the cases in which no shift in subject contexts is detected, but only supposed shifts in epistemic standards for knowledge attributions or denials, contextualism simply cannot give the right explanations for such cases. Indeed, the first strategy may yield certain results: examinees, when examined on or checked for their linguistic understanding of the predicate "(probabilistically) knows", would possibly apply all those different heuristics and manifest different "contents"; they would perhaps show that what they understand by it has something of their implicit mixing up the probability with representativeness, anchoring or some other "biased" strategy of reaching probabilistic conclusions. Therefore, any derived content of the predicate "(probabilistically) knows" would be surely contextually dependent. However, it could not count as relevant or adequate, since it would not be in accordance with the relevant invariant standard. The contextualist strategy is simply not so persuasive here as in the context of ordinary language considerations. Obviously, the point is extra-linguist in character: the truth values of the inferential knowledge attributions do not depend on any of the "lay" intuitions about the content of the relevant epistemic predicate. An analogous point holds for the possible strategy (2): the psychological examiners or, in that domain, knowledge attributors (as they were supposed to) had no practical interests concerning the quality of examinees' solutions in the case of the conjunction fallacy and in other similar cases. There was just nothing at stake for them.

Conclusion

Logical and probabilistic rules make *invariant* standards for the *attribution* of reasoning competence and respective inferential knowledge, as is abundantly confirmed by the experimental findings and their relevant interpretations by cognitive scientists. The presence or lack of knowledge in different contexts should be then explained by the intrinsic facts about the ways people perform their reasoning, i.e. about their current biases, cognitive histories or their natural and social environments, and not by the shifting standards for knowledge attributions. The *standards* for the *inferential knowledge* attributions – the epistemic standards of a particular

kind – remain the same across the contexts even when practices in different contexts may follow some other reasoning patterns (e.g. heuristics, domain specific rules or pragmatic reasoning schemas). The variations in the truth conditions and truth values of knowledge attributing or denying sentences in the cases of correct or fallacious reasoning described in the psychological literature were, therefore, sensitive only to the context of the reasoning subjects. In a generalized form, this thesis amounts to the claim that the attributions of inferential knowledge or ignorance (ignorance due to an inferential error) track the changes in the subject's context, given the fixed normative framework.

In the situations in which no shift in epistemic standards could be detected, but only shifts in the cognitive performance made by various subjects in various contexts, it seems that contextualism is not explanatorily better than invariantism. So, contrary to the contextualist thesis that "the truth conditions of knowledge-ascribing and knowledge-denying sentences (…) fluctuate in certain ways according to the context in which they are uttered" (DeRose 2002: 168), they fluctuate in dependence on the conditions in variable subject contexts.

This can be further reinforced by empirical findings, which, somewhat pessimistically, have already shown that *individual* subjects make fallacies in various contexts, from lay, everyday contexts to the expert ones, and that people in various communities and cultures make structurally similar fallacies and similar informal reasoning strategies. However, all this indicates the *existence* and *cognitive accessibility* (at least, to the experts) of *invariant* inferential and epistemic standards. That makes an optimistic argument in favor of invariantism.

References

Adler, J. E. and L. J. Rips (eds.) 2008. *Reasoning: Studies of Human Inference and Its Foundations* (Cambridge: Cambridge University Press).

Almog, J., J. Perry and H. Wettstein (eds.) 1989. *Themes from Kaplan*. (Oxford: Oxford University Press).

Cheng, P. and K. Holyoak 1985. "Pragmatic Reasoning Schemas", *Cognitive Psychology* 17, 391–416, in Adler and Rips (2008).

Cohen, S. 1998. "Contextualist Solutions to Epistemological Problems: Skepticism, Gettier, and the Lottery", *Australasian Journal of Philosophy* 76, 289–306.

——. 1999. "Contextualism, Skepticism, and the Structure of Reasons", *Philosophical Perspectives* 13, 57–89.

——. 2005. "Contextualism Defended", in Steup and Sosa (2005), 56–61.

Cosmides, L. 1989. "The Logic of Social exchange: Has natural selection shaped how humans reason?", *Cognition* 31, 197–276.

Cummins, D. 1996. "Evidence for the Innateness of Deontic Reasoning", *Mind and Language* 11, 160–90.

DeRose, K. 1992. "Contextualism and Knowledge Attributions", *Philosophy and Phenomenological Research* 52, 913–29.

——. 1995. "Solving the Skeptical Problem", *Philosophical Review* 104, 1–52.

——. 2002. "Assertion, Knowledge and Context", *Philosophical Review* 111, 167–203.

——. 2004. "Sosa, Safety, Sensitivity, and Skeptical Hypotheses", in J. Greco (ed.), *Ernest Sosa and His Critics* (Oxford: Blackwell), 22–41.

——. 2005. "The Ordinary Language Basis for Contextualism and the New Invariantism", *The Philosophical Quarterly* 55, 172–198.

——. 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context: Vol. 1* (Oxford: Clarendon Press).

Dretske, F. 1970. "Epistemic Operators", *Journal of Philosophy* 67, 1007–23.

——. 1971. "Conclusive Reasons", *Australasian Journal of Philosophy* 49, 1–22.

——. 2000. *Perception, Knowledge and Belief* (Oxford: Oxford University Press).

——. 2005. "The Case Against Closure", in Steup and Sosa (2005), 13–25.

Gigerenzer, G. and K. Hug (1992), "Domain-specific reasoning: Social contracts, cheating and perspective change", *Cognition* 43, 127–171

Gilovich, T, D. Griffin and D., Kahneman (eds.) 2002. *Heuristics and Biases: The Psychology of Intuitive Judgment* (Cambridge: Cambridge University Press).

Goldman, A. I. 1979. "What is justified belief?", in G. Pappas (ed.), *Justification and Knowledge* (Dordrecht: D. Reidel Publishing Company, 1979), 1–23.

——. 1986. *Epistemology and Cognition* (Cambridge, MA: Harvard University Press).

——. 1992. *Liaisons: Philosophy Meets the Cognitive and Social Sciences* (Cambridge, MA: MIT Press).

——. 2012. *Reliabilism and Contemporary Epistemology: Essays* (Oxford: Oxford University Press).

Hawthorne, J. 2004. *Knowledge and Lotteries* (New York and Oxford: Oxford University Press).

Heller, M. 1999. "The Proper Role for Contextualism in an Anti-Luck Epistemology", *Philosophical Perspectives 13: Epistemology*, 115–129.

Holland, J., K. Holyoak, R. Nisbett and P. Thagard 1986. *Induction: Processes of Inference, Learning, and Discovery* (Cambridge, MA: MIT Press).

Johnson-Laird, P. and R. Byrne 1991. *Deduction* (Psychology Press).

Kahneman, D., P. Slovic and A. Tversky 1982. *Judgment under Uncertanity*: *Heuristic and Biases* (Cambridge: Cambridge University Press).

Kaplan, D. 1989a. "Demonstratives", in J. Almog, J. Perry and H. Wettstein (1989), 481–563.

——. 1989b. "Afterthoughts", in J. Almog, J. Perry and H. Wettstein (1989), 565–614.

Kornblith, H. 1993. *Inductive Inference and Its Natural Ground*: *An Essay in Naturalistic Epistemology* (Cambridge: MIT Press).

Lewis, D. 1996. "Elusive Knowledge", *Australasian Journal of Philosophy* 74, 549–567.

——. 1999. *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press).

Schaffer, J. 2006. "The Irrelevance of the Subject: Against Subject-Sensitive Invariantism", *Philosophical Studies* 127, 87–107.

Sosa, E. 2000. "Skepticism and Contextualism", *Philosophical Issues* 10, 1–18.

Sperber, D., F. Cara and V. Girotto. 1995. "Relevance Theory explains the selection task", *Cognition* 57, 31–95.

Stanley, J. 2005. *Knowledge and Practical Interests* (New York and Oxford: Oxford University Press).

Steup, M. and E. Sosa (eds.) 2005. *Contemporary Debates in Epistemology* (Oxford: Blackwell).

Stich, S. 1990. *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation* (Cambridge, MA: The MIT Press).

Tversky, A. and D. Kahneman 1984. "Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment", *Psychological Review* 91, 293–315, in Gilovich, Griffin and D. Kahneman (2002), 19–48.

Unger, P. 1984/2002. *Philosophical Relativity* (Minneapolis: University of Minnesota Press).

Wason, P. 1966. "Reasoning", in B. Foss (ed.), *New Horizons in Psychology* (Middlesex, England: Penguin), 135–51.

Wason, P. and P. Johnson-Laird. 1972. *Psychology of Reasoning: Structure and Content*, (Cambridge: Harvard University Press).