The Generality Problem Naturalized

Olsson, Erik J

Published in:
Festschrift

2011

Citation for published version (APA):
Olsson, E. J. (2011). The Generality Problem Naturalized. In F. Svensson, & R. Sliwinski (Eds.), *Festschrift*

Total number of authors:
1

The Generality Problem Naturalized

Erik J. Olsson

Lund University

Abstract: In this paper, an attempt is made to get clearer on what, exactly, the notorious generality problem for reliabilism is all about. On a charitable reconstruction, one major issue is empirical, concerning as it does the frequency with which people converge in their assessments of reliability. Conee and Feldman's pessimistic view on this matter is contrasted with a more optimistic outlook that is taken to be supported by influential work in cognitive psychology. It is concluded that although much speaks in favor of the latter position, the dispute cannot be resolved by philosophical debating alone. Additional experimental work is also needed.

1. Introduction

Process reliabilism, or reliabilism for short, is the view that S knows that p if and only if (i) p is true, (ii) S believes that p, (iii) S's belief that p was acquired through a reliable process, and (iv) an appropriate anti-Gettier condition is satisfied. Reliabilism is sometimes advocated as a theory of epistemic justification, the main idea being that a person is justified in belief that p just in case her belief that p was formed via a reliable process, i.e., the process by means of which the belief was acquired was reliable.[1] For the purposes of the following discussion, there is no need to make a sharp distinction between these two brands of reliabilism. For definiteness, I will, initially, state much of the discussion in terms of knowledge.

The source of the longstanding *generality problem* for reliabilism is the observation that, because a process token is an unrepeatable causal sequence occurring at a particular time and place, it makes no good sense to ask whether a token process is reliable in itself. Rather, what can fundamentally be reliable or not are process *types*. A process token can still be said to be (un)reliable in a derivative sense if its associated process type is (un)reliable. For instance, the concrete process of Jones's coming to believe that he won the lottery on May 1, 2007, is itself neither reliable nor unreliable. However, given that its associated type is taken to be "belief formed through reading the local news paper", it is (probably) reliable.

The generality problem now arises because each token process can be classified as belonging to a great many different types, and it is not obvious how to single out one of these types rather than another as *the* unique associated type of the process in question. For example, the process leading up to Jones's belief could be classified narrowly as

---

[1] The process reliabilist account of knowledge was originally formulated by Ramsey (1931). For a modern (post-Gettier) account, see Goldman (1986). The process reliabilist theory of justification was first put forward in Goldman (1979).

belonging to the type whose sole member is Jones's coming to have his belief about the lottery, or, to take the other extreme, broadly as a belief formed through reading.

Furthermore, depending on what type is singled out as special, we may get different verdicts as regards the reliability of the process in question. Given the narrow classification in terms of the type whose sole member is the process producing Jones's belief, that process will, if the belief is true, be reliable. If it is seen instead as instantiating the general type "reading", it might be judged unreliable. Reading in general, irrespective of what is being read, is probably not reliable to an extent that suffices for knowledge.

These considerations reveal what appears to be an unacceptable lacuna in the reliabilist account of knowledge. From the reliabilist's perspective, whether a person knows or not will in many cases depend on whether the type of process producing the belief in question is reliable. And yet, as most commentators would agree, reliabilists have generally failed to clearly identify *the* type pertaining to a given token. In the absence of a principled account for how to select the relevant type, the reliabilist theory appears to be, in Conee and Feldman's words, "radically incomplete" (1998, p. 3).

Reliabilists have not been insensitive to the generality problem which was identified in Goldman (1979) and is portrayed as a serious issue for reliabilism in Goldman (1986) and, in particular, Goldman (2009). It is now considered to be a main challenge for a reliabilist theory by reliabilists and their critics alike.[2]

The problem has been stressed by Feldman (1985) and Conee and Feldman (1998). According to Feldman, solving the generality problem for reliabilism requires showing how to avoid the *single case problem* and the *no distinction problem*. The single case problem occurs when a process type is described so narrowly that only one instance of it ever occurs, and hence the type is either completely reliable or completely unreliable depending on whether the belief is true or false. The no distinction problem arises when beliefs of obviously different epistemic status are produced by tokens that are of the same (broad) relevant type. These two cases were illustrated above in our example with Jones's belief about the lottery.[3]

Conee and Feldman lay down three additional requirements for a solution to the generality problem. First, a solution must be "principled" in the sense of not being made on an *ad hoc* basis. Second, the rule must make reasonable epistemic classifications, by which is meant that the types identified must have a reliability that is plausibly correlated with the justificational status of the beliefs in question. Third, a solution must remain true to the spirit of the reliabilist approach and not characterize the relevant type of process in epistemic terms that are alien to reliabilist theorizing.

Conee and Feldman also provide a helpful classification of different solutions to the generality problem in terms of *common sense*, *scientific* or *contextually determined* types. As for the first approach, it is tempting to classify belief-forming process in terms of

---

[2] The generality objection can be found in many surveys of contemporary epistemology, e.g. Hetherington (1996), pp. 40-41, Lycan (1998), pp. 110-111, Plantinga (1993), p. 198, Pollock (1986), pp. 118-120, and Pollock and Cruz (1999), pp. 116-118. For a recent example, see Lemos (2007), pp. 92-94.

[3] The generality problem is usually stated as one of finding a *unique* relevant process type for each process token. Logically speaking, however, it would suffice to identify, for each token, a *class* of types that are either all reliable or all unreliable. Nevertheless, I will follow the mainstream and assume that the problem is to find a unique type for each token.

categories like "confused reasoning", "wishful thinking", or "hasty generalization". But, as Conee and Feldman note, reference to such commonsense types does not by itself solve the generality problem, for there are usually several common sense types that would fit a given token. Jones can be described as "reading a newspaper" or simply as "reading" where both descriptions plausibly pick out common sense types.

A second option is to identify an appropriate process type in scientific terms, i.e., by reference to "natural kinds". A proposal along these lines was made by Alston (1995), who wrote:

> With a process token, as with any other particular, any of its properties can be said to be correlated with a type to which it belongs … Even if it is true that you and I belong to indefinitely many classes, such as objects weighing more than ten pounds, objects that exist in the twentieth century; objects mentioned in this paper, etc. etc., it is still the case that membership in the class of human beings is fundamental for what we are in a way that those are not, just because it is the natural kind to which we belong. I shall suggest that something analogous is true of belief-forming processes – that there are fundamental considerations that mark out, for each such process token, a type that is something like its "natural kind" (p. 11).

Yet, there is, Conee and Feldman insist, a uniqueness problem facing this approach as well (pp. 10-11):

> What the natural kinds of belief-forming processes are is up for grabs, but every belief-forming process token is categorized in multiple ways by laws in each of several sciences. These all seem to be natural kinds of the process, according to current science. Reasonable candidates for natural kinds of a typical visual belief-forming process include electrochemical process, organic process, perceptual process, visual process, and facial-recognition process.

In his attempt to avoid this sort of criticism, Alston suggests narrowing down the different candidate types by appealing to the further criteria of *psychological realism* and *maximum specificity*. The idea is that the relevant type for a given process token corresponds to the maximally specific psychologically realistic natural kind to which the token belongs. In response, Conee and Feldman remark that "process reliability theories are supposed to appeal to much broader types" (p. 16).

Mark Heller (1995) offers a contextualist approach to the generality problem. Heller believes that there is a sense in which the generality problem, as it has been commonly understood, is unsolvable. For there is, Heller maintains, no fixed principle for selecting the relevant level of generality. In so far as the reliabilist is urged to supply such a fixed principle, the demand is unreasonable. Rather, what counts as correct varies from context to context. Conee and Feldman (1998) as well as Goldman (2008) find it doubtful whether an appeal to contextual factors suffices to narrow down a unique type for every context.

Having surveyed the different proposals that figure in the literature, only to find that they all fail to comply with at least one desideratum, Conee and Feldman conclude that "the prospects for a solution to the generality problem for process reliabilism are worse than bleak" (p. 5) and that "[i]n the absence of a brand new idea about relevant types, the problem looks insoluble" (p. 24), so that "process reliability theories of justification and knowledge look hopeless" (ibid.).

Since the appearance of Conee and Feldman's critical survey, a number of proposals have been put forward addressing the generality problem. Brandom (1998) advances a

theory that is contextualist in the sense that it focuses on third-person attribution of reliability featuring a reformulation of the generality problem in terms of defeasible inference. However, it is doubtful whether Brandom thereby really solves the problem or just reformulates it in inferentialist terms. Adler and Levin (2002) maintain that the generality problem is too general: if sound, it would apply also to processes mentioned in scientific explanations, which in their mind only shows that the generality problem is "illusory" to being with.[4] Comesaña (2006) argues, similarly, that the generality problem is no problem unique to reliabilist but that it is one that many justification-based theories share, Conee and Feldman's own evidentialist approach being a case in point. Kappel (2006) takes a stance that bears some resemblance to Heller's, arguing that there are no facts that determine a relevant type, but that this is not a fatal objection to the reliabilist theory but rather something that a proponent of such a theory can embrace.[5]

In the following I will question the assumption that the generality problem is a well-posed and well-understood problem. This will eventually lead to a proposed reconstruction of the problem in terms that should be attractive to practitioners of naturalized epistemology.

2. Generality reconsidered

First of all, what kind of theory is reliabilism supposed to be? What the reliabilist attempts to do is to provide necessary and sufficient conditions for the truth of "S knows that p", so much is clear. The reliabilist proposal, we recall, is that S knows that p if and only if (i) p is true, (ii) S believes that p, (iii) S's belief that p was acquired through a reliable process, and (iv) a suitable anti-Gettier clause is satisfied. Alternatively, reliabilism can, as we also saw, be thought of as spelling out necessary and sufficient conditions for justification in which case the suggestion is, roughly, the following: (RJ) S is justified in believing that p if and only if S's belief that p was acquired through a reliable process.

But, as Goldman noted in an early paper, "it is not enough for a theory [of justified belief] to state 'correct' necessary and sufficient conditions"; those conditions must also be "appropriately deep and revelatory" (Goldman, 1979, p. 92). A theory of justified belief must therefore "be couched at a suitably deep, general, or abstract level" (ibid.). While Goldman does not provide a systematic account of the desiderata that a theory of justified belief should satisfy in order to be acceptable, what he says in this regard can plausibly be subsumed under Carnap's general theory of philosophical explications (Carnap, 1950). According to that theory, a candidate explication of a pre-existing concept, the explicandum, should satisfy the following requirements to a sufficient degree (ibid., p. 7):

"1. The explicatum is to be *similar to the explicandum* in such a way that, in most cases in which the explicandum has so far been used, the explicatum can be used; however, close similarity is not required, and considerable differences are permitted.

2. The characterization of the explicatum, that is, the rules of its use (for instance, in the form of a definition), is to be given in an *exact* form, so as to introduce the explicatum into a well-connected system of scientific concepts.

---

[4] In their response to Adler and Levin, Feldman and Conee (2002) reiterate their view that the generality problem is real and remains unsolved.
[5] See also the exchange between Lepin (2007) and Christensen (2007).

3. The explicatum is to be a *fruitful* concept, that is, useful for the formulation of many universal statements (empirical laws in the case of a nonlogical concept, logical theorems in the case of a logical concept).

4. The explicatum should be as *simple* as possible; this means as simple as the more important requirements (1), (2) and (3) permit." [6] [7]

To take one of Carnap's own examples, in zoology the artificial concept piscis has come to replace the common sense concept of fish, although piscis is a narrower concept that excludes several kinds of animal which were subsumed under the concept fish, e.g. whales and seals. The reasons for the replacement is that the zoologists found that piscis is a much more fruitful concept than fish in the sense that it allows for the formulation of a greater number of interesting general truths. In the following I will reconstruct reliabilists as claiming, minimally, that their definition of knowledge or justification satisfies Carnap's four desiderata to a sufficiently high degree.

Suppose now that the critic is correct in thinking (A) that there are generally many different ways of classifying a belief formation process, and (B) that depending on how the process is classified we will get different verdicts as to whether or not the process was reliable. Let us finally also grant (C) that the reliabilist has so far been unable to devise a rule by means of which the right type can be identified in a given case. In what sense would these be disturbing facts for the reliabilist? Conee and Feldman's answer, of course, is that these facts together turn reliabilism into a "radically incomplete" epistemological theory. But what, exactly, is the reasoning leading up to this conclusion? Extracting a clear answer to this question from Conee and Feldman's work is surprisingly difficult. What follows is an attempt at reconstruction:

Perhaps Conee and Feldman are taking reliabilism to provide not merely a conceptual analysis, or explication, of knowledge or justification but also a method for deciding, in practice, whether someone knows or is justified. There are some indications in this direction in their 1998 paper, in which they write, for instance, that "[o]nly when a bearer of reliability has been identified does the theory have any implications about the justification of beliefs in particular cases" (p. 3). From this point of view, their criticism is relatively easy to appreciate: reliabilism is radically incomplete because it fails to give a method for identifying a process type "in particular cases". On the other hand, this interpretation of Conee and Feldman renders their criticism largely uninteresting for the simple reason that reliabilism is usually not advanced as a decision method along these lines, and in his 1985 paper Feldman explicitly distances himself from this understanding of reliabilism and the generality problem, writing "I do not assume that an acceptable solution to The Problem of Generality must provide a practical and useful method for identifying relevant types" (p. 173, footnote 6).

There is another, perhaps deeper, reason why this way of stating the generality problem is dubious. Critics of reliabilism universally present the generality problem as a challenge exclusively facing that theory. It is supposed to tell against reliabilism in a way that should make us more favorably disposed toward its main competitors, e.g., variants of JTB, evidentialism or virtue theory. This can be so only if there is no corresponding

---

[6] In the following I will adopt a more liberal understanding of fruitfulness according to which an explicatum can be fruitful in virtue of having general systematic benefits. Thus, an explicatum can be fruitful by avoiding problems that pertain to other ways of understanding the same explicandum, even if that explicatum is no more useful for the formulation of universal truths.

[7] See Maher (2007) for a recent defense of Carnap's method of explication.

problem arising for those other well-established analyses of knowledge. But if the complaint is that reliabilism does not provide a method for deciding in practice whether someone knows, then it is doubtful whether its critics have indeed succeeded in identifying a problem unique to reliabilism. After all, JTB analyses of knowledge, to take one example, do not provide such a method either because they typically do not provide a method for deciding, in particular cases, whether someone is justified in her belief. Providing, in Feldman's words, practical and useful methods of that kind is an interesting task in its own right but it is one that is orthogonal to the problem of providing good explications of knowledge or justification and one that concerns just about any account of knowledge or justification that comes to mind. Hence, the issue that will continue to be our concern here is to what extent there is a generality problem for reliabilism given the usual understanding of that theory as providing conceptual analyses, or explications, of some central epistemological concepts.[8]

Let us focus, for definiteness, on reliabilism as a theory of justification. Granted that the aim of an explication of justification is to give necessary and sufficient conditions for S's being justified in believing that p in a way that satisfies Carnap's desiderata to a sufficient degree, what would it mean for such an explication to be (radically) incomplete? Presumably that it failed to give either a necessary condition or a sufficient condition for justification. But (RJ) does state necessary and sufficient conditions for justification. It gives one condition that is claimed to be necessary as well as sufficient for S's being justified in believing that p, namely, that S's belief that p be acquired through a reliable process.

This leaves us with the alternative reconstruction of the criticism as targeting (RJ)'s ability to satisfy plausible desiderata on an explication of justification, such as those mentioned by Carnap. Let us first ask how well (RJ) fairs with regard to similarity to our everyday concept of justification. (RJ) can fail in this regard in two main ways: in being either too broad or too narrow. (RJ) is too broad if beliefs can be produced by reliable processes without being justified. It is too narrow if beliefs can be justified without being produced by reliable processes.

As for the latter case, suppose there are clear cases of justification, i.e. cases where we would all agree that a subject S is justified in believing that p. Imagine, for instance, that we are presented with a picture of Wilma standing in front of a tree in broad daylight. Wilma is, for all we can see, looking at the tree. There is nothing obscuring her view, we have no reason to think she is blind, and so on. Clearly we would then all assent to Wilma's being justified in believing that there is a tree in front of her. This is a clear case of justification. But is it also a clear case of belief acquisition through a reliable process? Given (A) – (C) one might be led to think that it is not. Given that there are many different ways to classify Wilma's belief formation process, leading possibly to widely different reliability assessments, and given that we have no explicit rule that we can

---

[8] The generality problem may also be thought of as the normative problem of how to resolve conflicting judgments concerning the proper categorizations of belief forming processes. Suppose X types process token t in one way whereas Y types that same token in another way, and that the effect is that X and Y differ in their reliability assessment. The problem is whether there is a rational way for X and Y to come to an agreement on how to type the process, e.g. by adopting a common type representing "suspension of judgment". Intriguing as this problem is, it arises in principle for all theories of justification, and not just for reliabilism. For any theory of justification, we may ask how it proposes to resolve conflicting judgments regarding justification. Hence this understanding of the generality problem is problematic for much the same reason as the "practical decision method" proposal is.

appeal to in the selection of one classification rather than another, it would seem that Wilma's is not a clear case of reliable belief formation. Rather, it will be, in a sense, indeterminate whether or not Wilma's belief was reliably formed.

This is indeed what certain theoretical considerations, namely (A) – (C), suggest. But why should we rest content with theory when we can study what happens in practice? Here is a sketch of a proposal for how this issue could be tested experimentally:

The proposal involves two groups of subjects confronted with, say, twenty episodes of ordinary life involving a person who comes to believe something for some reason or other (like Wilma above). For instance, the episodes could be presented to the subjects as film sequences.[9] Each subject in one of the groups is asked to state independently for each episode shown whether the person figuring in that episode is justified in her belief. Let us refer to this group as the *justification group*, or the J-group for short. Each subject in the other group is asked to state independently for each episode shown whether the person figuring in that episode acquired her belief in a reliable way. Let us refer to this group as the reliability group, or R-group for short. Obviously, the instructions for the subjects participating in this experiment should not contain any linguistic hints that may make them favor one classification rather than another.[10]

Assume now that this experiment has been carried out and that we got a good positive match between the reports of the J-group and the reports of the R-group. In other words, more or less the same the same episodes are described, by all or almost all members of the J-group, as involving justification and by all or almost all members of the R-group as involving reliable belief acquisition. That would be clear evidence in favor of the reliabilist theory of justification in the sense of supporting its faithfulness to our everyday concept of justification.

What would *disconfirm* that aspect of the reliabilist theory of justification? If we were to carry out the same experiments only to find there to be little or no positive correlation between judgments of justification and judgments of reliability that would be an unfortunate fact for the reliabilist theory, for that would indicate that Carnap's first desideratum is not satisfied after all.[11]

For the purposes of the following discussion we need to distinguish between two kinds of disconfirmation of (RJ). We would have a clear case of disconfirmation if, for instance, whenever all or almost all members of the J-group agreed that the person in the episode is justified, all or almost all members of the R-group agreed that the person did not acquire her belief in a reliable fashion. Let us refer to this as *category I disconfirmation*. But we could also encounter an outcome in which all or almost all members of the J-group agree in their justification assessments but in which there is strong dissonance among the members of the R-group concerning reliability. For instance, some members of the R-

---

[9] I am indebted to Martin L. Jönsson for the proposal to use film sequences for this purpose.

[10] We would like to use two groups in order to avoid various framing problems. If one group were to assess both justification and reliability, say in that order, the members' assessment of justification might conceivably influence their assessment of reliability.

[11] Similarly, a reliabilist theory of *knowledge* would be confirmed, with regard to faithfulness to ordinary language, if, in a significant number of cases, whenever all or almost all members of one group of subjects agrees that the person in the episode knows, all or almost all members of the reliability group agrees that the belief in question was reliably acquired. By contrast, if, in a significant number of cases, the members in the knowledge group agreed that the person knows, without the members of the second group reporting that the person's belief was reliably acquired, that would disconfirm the reliabilist theory of knowledge.

group may insist that the process was reliable, others may be equally convinced that it was unreliable and still others may fail to reach a verdict at all. Let us refer to this as *category II disconfirmation*.

To the best of my knowledge, experiments like the one just sketched have not yet been run, and so none of the parties in the generality dispute, including Conee and Feldman and myself, can claim knowledge about the outcome. This need not, however, prevent us from pondering on the likelihood of various alternative scenarios. As I construe them, Conee and Feldman report some rather strong views on this matter, to the effect that whereas the subjects in the J-group will, at least in paradigmatic cases, tend to give the same reports, the subjects in the R-group are likely to give unsystematic responses. In other words, Conee and Feldman predict that the experiment will result in category II disconfirmation of (RJ).

Why do Conee and Feldman predict that the subjects in the J-group will tend to give the same reports in paradigmatic cases? In our experiment, the episodes displayed to the participating subjects may be selected so as to exemplify what Conee and Feldman call "typical contexts in which knowledge and justification claims have clear truth value" (1997, p. 24). Conee and Feldman are committed to there being such contexts, and we must interpret them to think that people will give more or less identical reports when asked whether knowledge or justification is present in those contexts.

Conee and Feldman's argument for thinking that the subjects in the R-group will fail to give converging reports is more subtle. Here is a step-by-step reconstruction:

(CF1) People sometimes agree on typing and reliability but this happens only if the given type they converge upon is salient in the conversational context.

(CF2) The only way in which a type may be become salient in a conversational context is by means of linguistic presentation.

(CF3) Ordinarily, no type is linguistically presented.

(CF4) Hence, ordinarily, no type is salient and people will disagree on typing and reliability.[12]

This train of thought applies to our experiment because we were careful to assume that no type is linguistically indicated in the instructions to the participating subjects. Hence, Conee and Feldman offer the predication that the subjects in the R-group will be unable to converge on the same reliability assessment in concrete cases and that this failure is due to the lack of salient types. For instance, the number of R-group subjects reporting

---

[12] See Conee and Feldman (1998), p. 22-23, where they argue against Heller's claim that the relevant process type will contextually determined in many situations in which knowledge or justification is attributed: "There are some situations in which phrases referring to process tokens apparently work in the way Heller describes. For example, suppose Jones says, 'I have three ways to start my old jalopy: first, shifting into gear while rolling it down a hill; second, jump-starting it; and third, praying and then turning the key. Only the first two usually work.' Suppose that Jones then starts his car by jump-starting it. He remarks:
P. 'The process by which I just started my car is reliable'.
Here, Jones's explicit mention of the three types serves to limit drastically the types under consideration. The token mentioned in (P) is of one of those types only … In typical knowledge attributions, however, no contextual narrowing of candidate process types occurs … Ordinarily, no class of types of belief-forming processes will have been made contextually salient. And nothing else about typical contexts isolates any one type." See also Feldman and Conee (2002), p. 102, footnote 1.

Wilma's belief formation process to be reliable may be roughly the same as the number of subjects reporting it to be unreliable.

My own prediction differs significantly from Conee and Feldman's. On my prediction, both groups will report in a homogeneous manner, and there will be significant positive correlation between the reports of the two groups. For instance, most subjects in both groups will think of Wilma as simply 'seeing a tree', and this will make most subjects in the J-group report that Wilma is justified in her belief about the tree, and most subjects in the R-group report that the process by means of which Wilma formed her belief is reliable. I will proceed in the next section to add some substance to these claims.

But before doing that I would like to remark quickly on how well reliabilism fares regarding the remaining Carnapian explicatory dimensions: exactness, fruitfulness, and simplicity. Let us start with latter. The basic reliabilist analysis of knowledge in terms of truth, belief and reliable acquisition, is a strikingly simple proposal. Once the Gettier problem is taken into account, matters naturally become more complex, but there is no reason to think that reliabilism should fare worse than other analyses of knowledge in this respect.[13]

Is reliabilism a fruitful approach to knowledge and justification? This is a major issue in itself, which I cannot hope to do justice here. Historically, what has attracted epistemologists to reliabilism has often been the prospects of using it as an antidote to certain forms of skepticism. Knowledge is possible, according to reliabilism, if the external world exists more or less in the form we believe it does, and our beliefs about it are obtained through *de facto* reliable processes, even if we are not in position to experientially discriminate between various skeptical scenarios.[14] A further well-known systematic advantage of reliabilist theories of knowledge, over some more traditional accounts, is that it effectively blocks the regress problem. Finally, Goldman and others have argued recently that reliabilism can explain, in fairly natural ways, the distinctive value knowledge has over mere true belief. In some cases the explanation involves empirical law-like statements connecting reliabilist knowledge to other phenomena, such as future acquisition of true belief[15] or stability of true belief and successful action over time[16]. To a significant number of epistemologists, the theoretical and systematic virtues of reliabilism continue to present strong reasons in favor of that theory.

The issue of exactness is a delicate one. On first sight reliabilism seems to be as precise as any other epistemological theory that comes to mind. As Heller and others have insisted, the term "reliability" is as well-understood and unproblematic as one could hope in many everyday and scientific contexts. But the generality problem can be understood as shedding doubt precisely on that commonplace belief. On one rendering, the generality complaint is that "reliability" is too imprecise a term to be allowed to play a central role in epistemological theorizing. There is some truth to this proposal, for there is certainly vagueness concerning what *degree* of reliability must obtain in order for a process to count as reliable *tout court*. Reliabilists are often silent in this regard. But is there further, and deeper, imprecision concerning the typing of belief acquisition processes? And, if so, is that imprecision damaging to the theory? I am inclined to answer the first question in the positive and the second in the negative. There is imprecision as far as typing is

---

[13] Cf. Goldman (1986), Chapter 3, for a simple and elegant approach to the Gettier problem within a reliabilist framework.
[14] Cf. pp. 55-57 in Goldman (1986).
[15] Cf. the conditional probability solution in Goldman and Olsson (2009).
[16] Olsson (2007).

concerned, but that imprecision is not damaging to the theory, the reason being – as I shall now try to make likely – that the type is usually fixed, by contextual and cognitive factors, in ways that make us converge in our reliability assessments. In other words, in actual situations "reliability" is not much more imprecise than "you", "that car over there", "tomorrow" and other terms whose references are partly contextually determined.[17]

3. The case for convergence on reliability

One reason why we should expect people to give the same reliability verdicts in concrete cases comes from our ordinary life experience, for we often agree, it seems, about what is reliable and not, even in the complete absence of linguistic classificatory pointers. If someone comes forward claiming that this car (imagine the person pointing to a new Mercedes) is reliable, you would presumably tend to agree that it is. Similarly, we agree that the process by means of which the heart pumps blood through our veins is reliable (if only up to a point, unfortunately). Or suppose that we discover Karen sitting in the university library reading the Encyclopedia Britannica. Karen now raises her head saying, "The way I just learned the molecular structure of gold is reliable". Few of us would be inclined to object.

There is another reason for expecting that people should tend to converge on matters of reliability even if no relevant type-classification is made salient by linguistic means, for this is what influential empirically-based work in cognitive psychology indicates. In substantiating this claim I will draw on work on salience and so-called basic level effects in categorization.

Several psychological studies have been conducted on how people classify events and, in particular, what it is that makes some event categorizes more natural or salient. This work is relevant here because belief forming processes are kinds of events. Zacks and Tversky (2001) is a useful overview of the following and other advances in the psychological literature on event categorization. One line of research in this area indicates that what event type becomes salient may depend on the *time scale*. Consider "crossing the street", "walking to school", and "getting an education" as three different types corresponding to the same token event. According to the findings, event types

---

[17] A search on "reliable" in the ISI Web of Knowledge citation database resulted in a list of 17,955 scientific works having that word in the title (June 2010). The corresponding list for "known" and "justified" contained 13,629 and 2,499 items, respectively. The most common use of "justified" was in its ethical sense. These findings indicate that "reliable" is (a) a frequently used scientific term, and (b) more so than "known" and "justified". A corresponding comparison between "reliability" (59,506 hits), "justification" (4,049 hits), and "knowledge" (89,677 hits) resulted in the latter being the scientifically more established concept. But "reliability" was still found to be much more widely used than "justification". We recall Carnap's second desideratum for a good explication: "The characterization of the explicatum, that is, the rules of its use (for instance, in the form of a definition), is to be given in an *exact* form, so as to introduce the explicatum into a well-connected system of scientific concepts." (Carnap, 1950, p. 7). It seems that (RJ), the reliabilist explication of justified belief, is successful in this regard because it explicates a concept, justification, that is not very common in science in terms of another concept, reliability, that is in much more frequent scientific use and hence, we should expect, better connected to other scientific concepts. The extent to which the reliabilist explication of knowledge satisfies Carnap's second desideratum, taken literally, is less clear for the reason that "knowledge" is already a common scientific term, indeed more so than "reliability" itself. At least this is what this rudimentary bibliographical study suggests.

corresponding to a medium range time scale, such as "walking to school" are more likely to be salient than other more extreme types. This effect does not depend in any way on one or more types being made salient by linguistic means.

Here is an even more suggestive piece of evidence: According to one influential theory recognizing an event as an instance of a category consists of (A) matching it to a *schema* stored in memory and (B) matching features of the world to variables in the schema. Example of a schemata might be "X walks to Y", "X robs Y", and so on. Some schemata may be of evolutionary origin, others are developed in a social context through a learning process. The crucial point however is that an event type may become salient because it matches and activates a schema stored in memory. This phenomenon, too, occurs even if no type is singled out linguistically.

This is not the place to go into the details of how these phenomena transpire. What can be said with some confidence is that there is a general tendency in the relevant part of cognitive psychology to regard with suspicion any theory according to which a type can become salient in a given context only through explicit mentioning, Conee and Feldman's account being a case in point.

So far we have seen some evidence for thinking that non-linguistic factors can have salience-making powers. Such evidence does not by itself indicate that we should expect people to converge in their categorization of events of belief acquisition, for presumably any theory of salience, including one that highlights non-linguistic mechanisms, would have to allow for the possibility that several types can be salient in one and the same context, meaning that there will still be a need to explain how we could commonly converge on a *unique* type. Suppose, for example, that we have roughly the same structured process schemata stored in memory and that they serve to narrow down the psychologically real categories, in any given context, to just a few especially prominent ones. Little would prevent a given token process from potentially activating more than one such schema. For instance, Smith's seeing a tree could conceivably activate the following schemata all of which have some claim to psychological reality: *X perceives Y*, *X sees Y*, *X sees Y in bright sunlight*.

The forgoing remarks raise the question as to how people can come to agree on a natural *classificatory level* in taxonomies of salient types. The starting point from which will gradually approach an answer to that question will be object rather than event categorization. Consider the taxonomy in Figure 1.
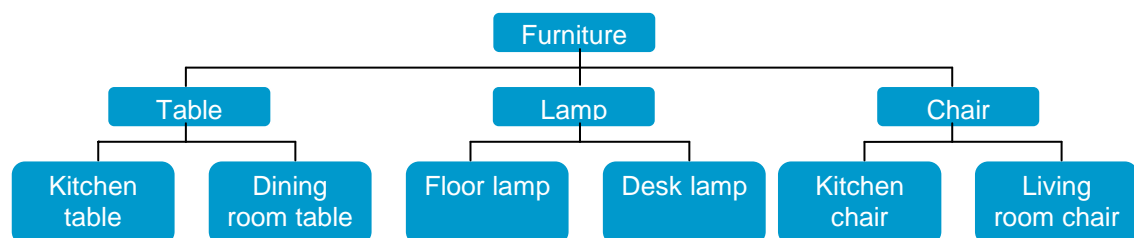
Figures 1: Taxonomy for furniture

Experimental studies of taxonomies as the one in Figure 1 have established that there is a level of categorization that is "basic" in the sense of being most natural. More specifically,

people will tend to think of things in terms of *Table*, *Lamp* and *Chairs* rather than in terms of superordinate (*Furniture*) or subordinate (*Kitchen chair* etc.) categories. Evidence for a basic level comes from a variety of experimental sources: basic level categories tend to be named when people are shown an object; in recognition tasks, people recognize basic level objects faster than either subordinates or superordinates; basic-level names generally have arisen earlier in the development of languages; basic categories are used earlier in the naming and other behavior of young children; finally, basic-level names tend to be shorter and more frequently used than names of superordinate or subordinate categories.[18]

The standard explanation of the basic level effect is based on the observation that categories tend to be associated with "attributes" which are things that can be inferred (in a defeasible sense) from membership in the category. For example, if a thing is classified as a *Chair*, we can infer that we can sit on it. If it classified as a *Lamp*, we can infer that it can be lit, and so on. Rosch et al (1976) list the following culturally significant attributes of furniture categories:

Furniture: no attributes

Chair: legs, seat, back, arms, comfortable, four legs, wood, holds people – you can sit on it

Kitchen chair: no additional

Living room chair: large, soft, cushion (additional attributes)

Thus, given that something is a chair, it can be inferred that the thing has legs, a seat, and so on. Given that something is a kitchen chair, the same inferences can be made and no new inferences are added. From the categorization of something simply as a piece of furniture, no useful conclusion can be drawn.

We can now explain the basic level effect by invoking two cognitive mechanisms assumed to be universally operative in human categorization. One mechanism favors economy of representation and the other usefulness (informativity) of categorization. We can conveniently speak of these mechanisms as representing two cognitive goals. The goal of economy requires that the classifier treat as many things as possible as "equivalent", i.e. as belonging to the same category. This goal will tend to favor broad categories. The goal of usefulness requires that the classifier seek a maximally useful classification of a given thing. A classification is useful to the extent that interesting attributes of the object can be inferred, again defeasibly, from category membership alone. This goal will tend to favor narrow categories.[19] The challenge facing a classifying subject is to make a reasonable trade-off between these two goals, i.e., to choose a type that is informative and yet cognitively parsimonious.

Rosch and her associates proposed a specific way of striking the right balance: the basic classificatory level is "the most inclusive level in a taxonomy at which a cluster of attributes, believed to be common to the class named, would be listed" (Rosch et al, 1976, pp. 435-436, original emphasis removed). Let us refer to this rule for identifying the basic level as the *Rosch rule*. For instance, the *Chair* level is the basic level because it involves a cluster of attributes which are not listed at higher levels in the taxonomy. The *Kitchen*

---

[18] Cf. Rogers and Patterson (2007), p. 451: "… basic level effects are among the best known and most robust phenomena in the study of human categorization."
[19] See Rosch et al (1976), pp. 383-385, for details.

Chair level is not basic because, although it too involves a cluster of listed attributes, those attributes are listed at the higher *Chair* level as well.[20] [21]

I will now propose to treat the classification of belief forming processes in an analogous fashion. The proposal is to use the Rosch rule to determine the type for a given process token just as we did for concrete objects. This is in line with the observation of basic level effects in the categorization of events (of which belief formation processes are a special case) and the claim made by Rosch and her colleagues that the principle of category formation is universally valid.[22] The two cognitive goals previously alluded to are operative in the categorization of belief formation processes, so that our natural inclination is to obtain a categorization that is both economical and informative, just as before. The most basic level in a process taxonomy is the most inclusive level at which a cluster of attributes, believed to be common to the class named, would be listed.

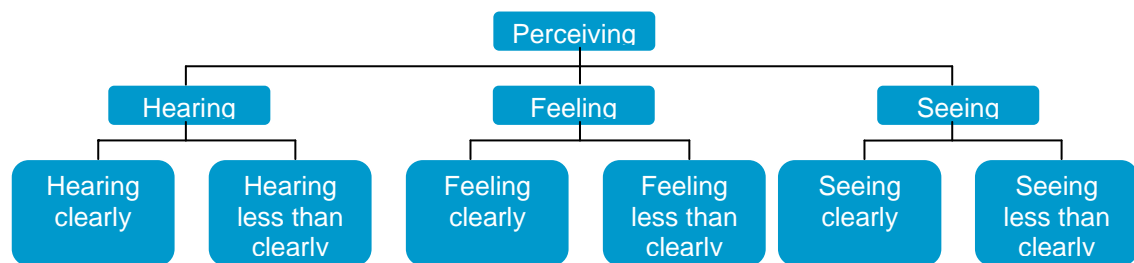The following example illustrates how this is supposed to work.

Figure 2: Hypothetical taxonomy for belief formation

I hypothesize that the following attributes would be considered to be of potential importance for most people:[23]

Perceiving: no attributes

Seeing: reliable (for most practical purposes), believing, proximity to object, nothing (opaque) obstructing the view[24]

Seeing clearly: very reliable (for scientific/legal purposes)

Seeing less than clearly: no additional[25]

---

[20] The Rosch rule suffers from imprecision concerning how many attributes it takes to form a "cluster". This potential shortcoming turns out to be inconsequential in the present context, as we will focus on a special case of the rule in which the notion of an attribute cluster plays no role.

[21] Variants of the explanatory model originally proposed by Rosch and colleagues have continuously been employed in order to explain phenomena connected with categorization and non-logical reasoning. See Jönsson and Hampton (2006, 2008) for recent examples.

[22] Rosch et al (1976) wrote that "the principle of category formation … is claimed to be universal" (p. 435) and not valid merely for concrete objects. See also Mervin and Rosch (1981), p. 93. Zacks and Tversky (2001) discuss some work on basic level effects in event categorization.

[23] I am not aware of any actual experiments designed to test the basic level theory for the special case of belief formation processes. Nor am I, for that matter, aware of any evidence suggesting that such processes should constitute an exception to the general rule that classification of events conforms to that theory.

[24] "Seeing" is here and elsewhere taken in its non-veridical sense.

The basic level is given by the most inclusive level classification of visual perception at which a cluster of attributes are listed, which in this case is the level of *Seeing*.

The account we have given of type convergence also predicts convergence in matters of reliability. For the reasons given, people will tend to categorize belief formation processes in terms of *Seeing*, *Hearing*, and so on. They will then tend to infer the attributes associated without those categories, one of those attributes being "reliable (for most practical purposes)".

Let us return to the previous example of Wilma for a concrete illustration. I claimed that most of the subjects in the R-group would think of her as simple "seeing a tree", a process they will report to be reliable (at least for everyday purposes). The reason is that *Seeing* is a basic level event category and *Tree* a basic level object category.[26] Either category allows many inferences to be drawn to potentially useful attributes. Combining the categories into "seeing a tree" is a cognitively attractive move because it results in a description that is compact yet loaded with information content.

Someone might object to our explanation of typing convergence on the following basis: "Fine, maybe this succeeds in showing, at least in outline, how we can agree on a process classification and reliability in some cases, but these cases do not seem to correspond to those that epistemologist take special interest in. Some of the 'attributes' of *Seeing* – like 'proximity to the object' – are not really epistemically interesting, which raises the question why we should base our classifications on them. In epistemic contexts, what we care about is simply the reliability of the process, and that is all there is to it. What is needed is an argument as to why we tend to agree on classification/reliability when the reliability only is at stake. The Rosch-based theory presented here doesn't seem to do that job."

It is true that what has been said so far assumes that we are interested in a *multi-purpose classification* where many attributes are (potentially) important.[27] What the objector points out is that the original classificatory problem can also be understood as one of *single purpose classification*. On this interpretation, there is only one attribute/property that we care about: reliability. What happens in this case? Once more, the original Rosch rule for multi-purpose classification states that the basic level of classification is the most inclusive level in a taxonomy at which a cluster of attributes, believed to be common to the class named, would be listed. This suggests that, in the single purpose case, the basic level of classification is the most inclusive level in a taxonomy at which the attribute of interest would be listed.

There is however an immediate problem with this proposal. Suppose the taxonomy consists of the superordinate *Being deceived* and the subordinates *Being deceived intentionally* and *Being deceived unintentionally*, and that the token process in question is

---

[25] Alternatively, "not very reliable (for scientific/legal purposes)" could be added as an attribute of the category *Seeing less than clearly*.

[26] Rosch et al (1976) found empirical support for the claim that *Tree* is a basic level object category (pp. 390-391).

[27] Cf. Corter and Gluck (1992), p. 293: "Our basic assumption is that there is functional value for a person to have accurate information about the features of things. For example, an organism searching for food needs to know whether a given plant part is poisonous, nutritive, sweet, tough, and so on. Some features of instances may be useful only indirectly – for example, to generate tests to confirm tentative identifications. However, because a person will experience a variety of need states and goals across time, generally, there is value for the person to have accurate information about all the features of instances."

a case of, say, intentional deception. Since "reliable" is not an attribute of any category to which the token process belongs, there is no level in the taxonomy at which "reliable" is listed. Hence, no type is singled out by the rule.

The problem is quickly solved by adding that the relevant taxonomic levels are such that either the attribute itself *or its negation* is listed. In other words, *the basic level of classification, in the single purpose case, is the most inclusive level in a taxonomy at which the attribute or its negation would be listed*. To illustrate, the following attribute list would be relevant if the task is single purpose categorization with respect to reliability (for most practical purposes):

Perceiving: no attributes

Seeing reliable (for most practical purposes)

Seeing clearly: no additional

Seeing less than clearly: no additional

If, as we assume, the process was a case not only of *Seeing* but also of *Seeing clearly*, both these categories satisfy the condition that "reliable (for most practical purposes)" is listed. However, *Seeing* is the most inclusive category satisfying that condition, which is why it is singled out by the single purpose Rosch rule as the type of the process in question.[28]

So far we have been looking at taxonomies that are *monotonic* in the sense that subordinate categories "inherit" all attributes from their superordinates.[29] In the furniture case, for instance, the *Chair* attribute "has a seat" was supposed to be inherited by the subordinate categories of *Kitchen Chair* and *Living Room Chair*. Likewise, the reliability attribute of *Seeing* was assumed to be inherited by its subordinate categories. Monotonic taxonomies admit a highly compressed cognitive representation which allows us to store attributes only once without having to repeat them for subordinate categories. For that reason, we should expect the human mind to favor storing information in monotonic taxonomies if that is an option. We have provided an account of reliability convergence in monotonic taxonomies. But what about the non-monotonic case?

Suppose, to be specific, that John's process leading up to his belief that the person over there is Susan instantiates the type *Seeing*. This does not prevent it from also instantiating the type *Occluded seeing*, by which is meant that the target of vision is at least partly occluded. If classified as *Seeing*, the process can be inferred to be reliable (for everyday purposes), whereas *Occluded seeing* will be associated with the attribute "unreliable". The matter doesn't end here: suppose that in the case we are discussing, the occlusion only pertains to the target's periphery, and furthermore that the occlusion only blocks S's view of Susan's body, without blocking the view of her face. Let us call this a case of *Insignificantly occluded seeing*. If the process is typed as *Insignificantly occluded seeing*, we may once more infer reliability. And so things could continue.[30]

---

[28] The result of single purpose classification might diverge from the result of multi-purpose classification, which is in this case *Seeing*. Take as the single-purpose attribute that of being very reliable. In Figure 2, "very reliable" appears only at the level of *Seeing clearly*, which is therefore the single purpose basic level category.

[29] This is the term used by Corter and Gluck (1992), p. 296.

[30] Thanks to Alvin Goldman for providing this example (personal communication). Brandom (1998) describes analogously unending switches in reliability as one moves from one process category to another.
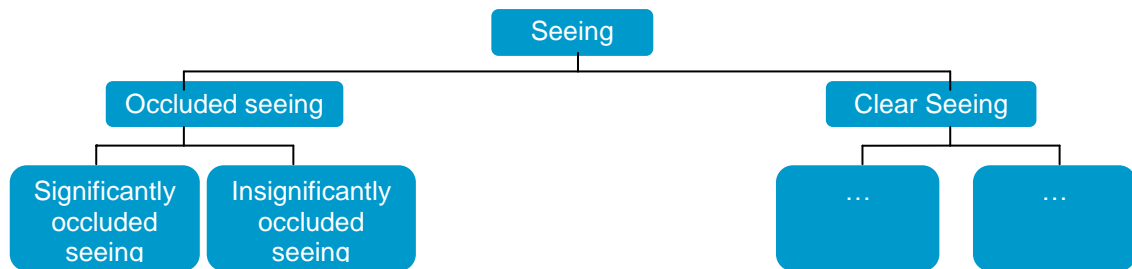
Figure 3: A non-monotonic taxonomy

If we focus on the single purpose classificatory task, the attribute of reliability/unreliability would be assigned as follows:

Seeing: reliable

Occluded seeing: unreliable

Insignificantly occluded seeing: reliable

Significantly occluded seeing: unreliable

One can imagine how a taxonomy like the one in Figure 3 could be generated "on the fly" in the context of a dialogue:

John: I know Susan is over there. I see her.

Mary: How can you be sure given the occlusion?

John: True, but I see her face so the occlusion is insignificant.[31]

Logically speaking, nothing prevents ever more fine-grained types from being proposed endlessly, but that would surely be quite unrealistic in practice. Rather, we would expect the process of generating types on the fly soon to come to an end. At that point, there are a finite number of types "on the table" forming a taxonomy of types. Is there a basic level even if the taxonomy is non-monotonic and, if so, what characterizes that level?

The Rosch rule for single purpose categorization states that the basic level is the most inclusive level in the taxonomy at which reliability or unreliability would be listed. If the occlusion was in fact insignificant, this means that the Rosch rule singles out the category of *Seeing* as the basic level category, from which we may infer that the process was, as we would have expected it to be, reliable. The problem is that the Rosch rule gives exactly the same result if we assume instead that the process was one of significantly occluded seeing. In that case, too, we get the result that *Seeing* is the basic level category so that the process was reliable, which is simply incorrect. While the Rosch rule for

---

[31] Non-monotonic effects also arise in scientific taxonomies. In biology, the category *Whale* is subordinate to the category *Mammal* and yet the mammal attribute "lives on land" is not inherited by the whale category. Another classic example involves penguins, i.e., birds that, unlike most, do not fly.

single purpose categorization is plausible for monotonic taxonomies, it fails to account for basic level effects in the presence of non-monotonicity.

An alternative suggestion comes to mind: that we should focus initially on categories to which the process belongs that cannot be further subdivided in ways that would make a difference to what can be inferred about the reliability or unreliability of the process in question; and that we should choose, among the candidates that remain, the category that is most inclusive. Suppose again that the occlusion was insignificant, which means that the token process belongs to three categories: *Seeing*, *Occluded seeing* and *Insignificantly occluded seeing*. Only the latter cannot be further subdivided in a way that is relevant to reliability, which means that it will be singled out as the process type, enabling us to infer that the process was reliable. By the same token, if the occlusion was significant, the process will be typed as *Significantly occluded seeing*, from which we may infer that it was unreliable.

This rule gives the correct verdict concerning reliability, based on a plausible classification of the process in question. Nevertheless, I have one, perhaps minor, concern regarding the prospects of justifying it within a Rosch style theoretical framework. As we recall, two cognitive goals are taken, within that framework, to be universally operative in human cognition: the goal of economy and the goal of usefulness. It could be agued that, while the alternative rule may be justifiable from the point of view of usefulness, it is suboptimal with respect to economy. Let us focus on the latter point. A process that is a case of insignificantly occluded seeing is assigned the type *Insignificantly occluded seeing* by the statistically inspired rule. Based on that classification, it can be inferred that the process was reliable. But the same inference could have been made given a coarser and more economical typing of the process as *Seeing*. Similarly, a process that is a case of significantly occluded seeing is assigned the type *Significantly occluded seeing*, which allows an inference to unreliability. Once more, the same inference could also have been drawn by more economical means by typing the process as *Occluded seeing*.

Here is an amended rule that improves on the proposal just made with regard to the economy aspect. Let us say that the single purpose category for token process t for taxonomy T (with respect to reliability) is the most inclusive T-category C such that (i) t belongs to C and (ii) C is associated with the same reliability attribute (reliable/unreliable) as is the most specific T-category to which t belongs. In other words, we first check what can be inferred regarding reliability from the most specific categorization of a given token process in the taxonomy and then seek the most inclusive superordinate from which that same inference can be drawn. As can easily be checked, the amended rule will fare better than the statistically inspired rule with respect to economy because it will type insignificantly occluded seeing as *Seeing* and significantly occluded seeing as *Occluded seeing*.

This completes my motivation for the claim that people will tend to converge in their assessments of reliability. My aim has not been to provide a complete psychological story for why convergence should be expected. I believe however that enough has been said to shift the burden of proof to Conee and Feldman, who, as we recall, predict massive disagreement on reliability in the absence of linguistic guidance, a predication for which they offer no evidence beside their own intuitions. Ultimately, the matter will have to be decided through further empirical investigation.

4. Generality and the context of inquiry

One more objection can be anticipated. To return to our hypothetical experiment, the account given in the previous section supports the claim that the subjects in the R-group will concur with respect to process typing and reliability, at least within reasonable limits, and the types they will concur on will typically be common sense types, like *Seeing*, *Hearing*, and so on. Conee and Feldman, however, believe that such types are deeply problematic in an epistemological context. Take, for instance, the common sense type *Brief and hasty scanning*:

> Sometimes, on the basis of brief and hasty scanning we can get extremely well justified beliefs, as when we see in a glance that there is a tree in the backyard. Other times brief and hasty scanning does not yield a justified belief, as when the belief concerns exactly how many leaves there are on the tree. Simple common sense classifications are thus too broad to make the right epistemic distinctions among beliefs (p. 7).

The thought here is that while brief scanning can sometimes result in a justified belief and other times result in an unjustified belief, such scanning will have to be either reliable in all contexts or unreliable in all contexts. Therefore, brief scanning will be "too broad to make the right epistemic distinctions" (ibid.).

It seems true that brief scanning can sometimes result in a justified belief and other times result in an unjustified belief. But it seems equally true that brief scanning is sometimes a reliable process and sometimes an unreliable process. Roughly, brief scanning is a reliable process if the belief in question concerns an easily detectable object or feature of an object, and unreliable otherwise. Suppose for instance that the subjects participating in our experiment are shown a film sequence of Wilma taking a quick look through her window and writing on a piece of paper "There is a tree in my backyard". We now ask the subjects in the R-group how Wilma came to believe that there is a tree in the backyard, whereupon they may very well respond that she briefly scanned the backyard. The subjects may then add that the process was reliable (because the object was easily detectable). Imagine instead that the same film sequence is shown but it now ends with Wilma writing on a piece of paper "The tree in my backyard has 55,678 leaves". The subject may very well classify the process similarly as featuring Wilma taking a quick look at the backyard. However, this time it seems perfectly possible that the subjects will report that the process was unreliable (because the object feature in question was not easily detectable).

Thus the context of inquiry will influence not only our judgment as to whether a belief is justified, but also our judgment as to whether the process leading up to it was reliable. We can enter different contexts of inquiry by posing different question to the subjects participating in our hypothetical experiment. If we ask them how Wilma came to believe there is a tree in the backyard, then this will influence the subject in a certain direction. If we instead ask them how Wilma came to believe that the tree has 55,678 leaves, this will influence the subjects in another direction. A plausible psychological theory of salience should allow for the context of inquiry to influence which attributes of categories become salient. The bottom line is that all this is fine so long as both groups – the J-group and the R-group – are confronted with the same question *mutuatis mutandis*. In other words, either we query both groups about the belief about the tree or we query both groups about the belief about the exact number of leaves on the tree. Reliabilism predicts that we will get positive correlations between the answers to the groups *within one and the same context of inquiry*. This is intuitively the most charitable understanding of the reliabilist theory, and it also follows from the proposal that the reliabilism should be seen as

advancing explications in Carnap's sense. It is a requirement on an explication, we recall, that the explicatum be similar to the explicandum in such a way that "in most cases in which the explicandum has so far been used, the explicatum can be used" (Carnap, 1950, p. 7). Hence, in the first scenario, featuring the belief about the presence of the tree in the backyard, reliabilism predicts that, in most cases in which members of the J-group say that the belief is justified, members of the R-group will say that it was reliably acquired. The same goes for the second scenario, featuring the belief about the number of leaves. Reliabilism does not predict that we get positive correlations across different contexts of inquiry. It does not predict, for instance, that the answer given by the J-group in the first scenario will match the answer given by the R-group in the second scenario.

5. Conclusion

In this paper, I have questioned the assumption that the generality problem is a well-understood issue. This led to the discovery that underlying the generality debate there is an empirical issue concerning what we can expect regarding peoples' natural tendencies to type processes in similar ways and, based upon the common type assignment, offer the same reliability judgments. Conee and Feldman were seen to advocate, on purely philosophical grounds, a linguistic theory of type salience on the basis of which it can be predicted that people will rarely converge in their reliability verdicts. This hypothesis was contrasted with what seems to be the received view in cognitive psychology, namely, that non-linguistic factors, too, can contribute to salience. While this better underpinned account of salience does not by itself entail that there will be convergence in judgments of reliability, the combination of that theory with the standard Rosch basic level theory should make us expect convergence in many cases, or so it was argued.

A final remark: several authors have argued that the generality problem is in fact not a serious issue for reliabilism, e.g., because the problem "cuts across the board" (Adler and Levin, 2002, p. 97), because it affects competing theories as well (Comesaña, 2006), or because "unreasonable demands [are] placed upon the reliabilist" (Heller, 1995, pp 502-503). By contrast, the message here is that the generality problem can be understood, more charitably, as raising some real and, apparently, still largely open questions in the intersection between epistemology and psychology regarding (a) how people actually type belief formation processes, and (b) to what extend there is widespread convergence on the types thus assigned and on the corresponding reliability judgments. Hence, one merit of the present account, apart from the obvious prospects it offers for unifying work in epistemology and cognitive psychology, is that it seems to do greater justice than the alternative accounts just mentioned to the fact that leading epistemologists like Richard Feldman and Alvin I. Goldman have identified the generality problem as a significant concern for the reliabilist theory.

References

Adler, J., and Levin, M. (2002), "Is the Generality Problem too General?", *Philosophy and Phenomenological Research* 65 (1): 87-97.

Alston, W. (1995), "How to Think about Reliability", *Philosophical Topics*, Spring: 1-29.

Beebe, J. R. (2004), "The Generality Problem, Statistical Relevance and the Tri-Level Hypothesis", *Nous* 38: 1, 177-195.

Brandom, R. B. (1998), "Insights and Blindspots of Reliabilism", *The Monist* 81 (3): 371-392.

Carnap, R. (1950), *Logical Foundations of Probability*, Chicago: Chicago University Press.

Christensen, D. (2007), "Three Questions about Leplin's Reliabilism", *Philosophical Studies* 134 (1): 43-50.

Comesaña, J. (2006), "A Well-founded Solution to the Generality Problem", *Philosophical Studies* 129: 27-47.

Conee, E. and Feldman, R. (1998), "The Generality Problem for Reliabilism", *Philosophical Studies* 89: 1-29.

Corter, J. E., and Gluck, M. A. (1992), "Explaining Basic Categories: Feature Predictability and Information", *Psychological Bulletin* 111 (2): 291-303.

Dupré, J. (1993), *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*, Harvard University Press.

Feldman, R. (1985), "Reliability and Justification", *The Monist* 68: 159-174.

Feldman, R., and Conee, E. (2002), "Typing Problems", *Philosophy and Phenomenological Research* 65 (1): 98-105.

Goldman, A. I. (1979), "What is Justified Belief?", in Pappas, G. (ed.), *Justification and Knowledge*, Dordrecht: D. Reidel.

Goldman, A. I. (1986), *Epistemology and Cognition*, Cambridge MA: Harvard University Press.

Goldman, A. I. (1992), "Reliabilism" in *A Companion to Epistemology*, J. Dancy and E. Sosa (Eds.), Cambridge, MA: Blackwell.

Goldman, A. I., "Reliabilism", *The Stanford Encyclopedia of Philosophy* (Summer 2008 Edition), Eward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2003/entries/davidson/>.

Goldman, A. I., and Olsson, E. J. (2009), "Reliabilism and the Value of Knowledge", in Haddock, A., Millar, A. and, Pritchard, D. H. (eds.), *Epistemic Value*, Oxford University Press.

Heller, M. (1995), "The Simple Solution to the Problem of Generality", *Noûs* 29 (4): 501-515.

Hetherington, S. C. (1996), *Knowledge Puzzles: An Introduction to Epistemology*, Boulder, CO: Westview.

Jönsson, M. L., and Hampton, J. A. (2006), "On Prototypes as Defaults", *Cognition* 106 (2): 913-923.

Jönsson, M. L., and Hampton, J. A. (2008), "The Inverse Conjunction Fallacy", *Journal of Memory and Language* 55 (3): 317-334.

Kappel, K. (2006), "A Diagnosis and Resolution to the Generality Problem", *Philosophical Studies* 127: 525-560.

Lemos, N. (2007), *An Introduction to the Theory of Knowledge*, New York: Cambridge University Press.

Leplin, J. (2007), "In Defense of Reliabilism", *Philosophical Studies* 134 (1), 2007: 31-42.

Lycan, W. G. (1988), *Judgment and Justification*, New York: Cambridge University Press.

Maher, P. (2007), "Explication Defended", Studia Logica 86 (2): 331-341.

Olsson, E. J. (2007), "Reliabilism, Stability, and the Value of Knowledge", *American Philosophical Quarterly*, vol. 44 (4): 343-355, 2007.

Plantinga, A. (1993), *Warrant: The Current Debate*, Oxford: Oxford University Press.

Pollock, J. L. (1986), *Contemporary Theories of Knowledge*, Totowa, NJ: Rowman and Littlefield.

Pollock, J. L., and Cruz, J. (1999), *Contemporary Theories of Knowledge*, 2nd edition, Lanham, MD: Rowman and Littlefield.

Ramsey, F. P. (1931), *The Foundations of Mathematics and other Logical Essays*, Braithwaite, R. B. (Ed.), London: Routledge and Kegan Paul.

Rogers, T. T., and Patterson, K. (2007), "Object Categorization: Reversals and Explanations of the Basic Level Advantage", *Journal of Experimental Psychology: General* 136 (3): 451-469.

Rosch, E,. Mervis, C. B, Gray, W. D, Johnson, D. M, and Boyes-Braem, P. (1976), "Basic Objects in Natural Categories", *Cognitive Psychology* 8: 382-439.

Wunderlich, M. (2003), "Vector Reliability: A New Approach to Epistemic Justification", *Synthese* 136: 237-262.

Zacks, J. M., and Tversky, B. (2001), "Event Structure in Perception and Conception", *Psychological Bulletin* 127 (1): 3-21.