Episteme, 10, 3 (2013) 241–268 © Cambridge University Press
doi:10.1017/epi.2013.27

A RELIABILISM BUILT ON COGNITIVE CONVERGENCE: AN EMPIRICALLY GROUNDED SOLUTION TO THE GENERALITY PROBLEM

MARTIN L. JÖNSSON
martin.jonsson@fil.lu.se

ABSTRACT

Process-reliabilist analyses of justification and knowledge face the generality problem. Recent discussion of this problem turns on certain untested empirical assumptions that this paper investigates. Three experiments are reported: two are free-naming studies that support the existence of a basic level in the previously unexplored domain of names for belief-forming processes; the third demonstrates that reliability judgments for the basic-level belief-forming process types are very strongly correlated with the corresponding justification and knowledge judgments. I argue that these results lend support to process-reliabilism; I then go on to explore several different ways to make use of them in solutions to the generality problem.

According to a simple process-reliabilist theory of knowledge, a person *A* knows that *p* if and only if: i) *A* believes that *p*, ii) *p* is true, and iii) *A*'s belief that *p* was obtained by a reliable process.[1] Similarly, according to a simple reliabilist theory of *justification*, a belief is justified if and only if it was obtained by a reliable process. On both accounts, a process is assumed to be reliable if it tends to produce true beliefs; this, in turn, might be given an interpretation in terms of actual frequencies or frequencies across certain possible situations.

It is generally assumed that both types of reliabilism face the *generality problem*: while every belief is produced by a dateable *token* process, reliability (which is a tendency) is a property of a *type* of process.[2] Since each token-process belongs to numerous process-types that differ in reliability, reliabilism – as stated above – is under-specified in a way that makes its predictions, in a large class of cases, entirely unclear. In the words of two of the main discussants of the problem:

> The token event. . .of [Smith seeing a maple tree through a window] is an instance of the following types, among others: visually initiated belief-forming process, process of retinal image of such-and-such specific characteristics leading to a belief that there is a maple tree nearby, process of relying on a leaf shape to form a tree-classifying judgment, perceptual process of classifying by

1 There are different kinds of reliabilism. The kind that I will exclusively concern myself with in this paper is process-reliabilism, a variety of reliabilism originally due to Goldman (1979). It should be noted that Goldman's original proposal was to analyze justification (and not knowledge) in terms of reliable processes. See, for instance, Goldman (1976) for a form of counterfactual reliabilism for knowledge that is not a form of process-reliabilism.

2 The generality problem was first discussed by Goldman (1979:12). See Conee and Feldman (1998) for a good overview of attempts to solve the problem.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

> species a tree located behind a solid obstruction, etc. The number of types is unlimited. They are as numerous as the properties had by the belief-forming process. Thus, process reliability theories confront the question of which type must be reliable for the resulting belief to be justified. It is clear that the answer to this question will significantly affect the implications of the theory.
>
> (Conee and Feldman 1998: 2)

Although the quote clearly draws attention to a *prima facie* problem for reliabilism, a quick review of the different process-types suggested in the quotation reveals one of them to be considerably more natural to use than the others when reporting the relevant token-process. Most people would probably prefer 'Smith *saw* a maple tree through the window' over one of the alternatives. If this is true, this fact is a special case of a very general form of cognitive and linguistic convergence: people in the same kind of situation tend to type tokens at the same level of abstraction. That this form of convergence actually obtains in general is one of the core discoveries of the cognitive psychologists concerned with *basic-level concepts*. As a striking illustration of the joint agenda of the reliabilists and these cognitive psychologists, compare the following quote – from an overview of the research on basic-level concepts – with the previous one.

> The people, objects, and events that we encounter every day do not each fit into a single category. Wilbur might simultaneously be a bulldog, a dog, a mammal, and an animal...People fit into many categories such as being a woman, a reporter, a political conservative, a Yankees fan, a New Yorker, an African-American, a cousin and so on...The same is true for events, as the same event might be thought of as seeing a comedy, going to the theater, spending an evening out, getting entertainment, going on a date, or (hopefully not) a fiasco....The question arises... How do we decide which of these categories to use on any occasion?
>
> (Murphy 2002: 199)

Since the mid-seventies, basic-level research has been concerned with answering this question by developing theories about what is special about basic-level words and concepts, and by demonstrating that there is, quite generally, a basic level in thought and language to which some concepts and words belong.

In a recent discussion of the generality problem, Erik Olsson (forthcoming) is the first to draw attention to the basic-level literature as an untapped resource for solving the generality problem.[3] In light of the amassed empirical data from other domains, Olsson argues that it is likely that people normally also type process-tokens in a homogenous manner, and that they thus converge on what we may call 'intuitive process-types'. He also suggests that the reliability of these types will be able to predict the corresponding justification and knowledge judgments. On his view, given that intuitive process-types exist, process-reliabilism can avoid the underspecification that the generality problem draws attention to. However, the existence of intuitive process-types is not uncontroversial: Feldman and Conee (2002) maintain – in reply to claims made by Adler and Levin

3 Olsson's approach to justification and knowledge is a descriptive attempt to capture the folk concepts of these notions (cf. Goldman 1992), and the brand of reliabilism he has on offer is thus something that should be evaluated in light of the intuitions people have concerning the normal attribution of these notions to themselves and others.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

(2002) – that, in the absence of an explicit linguistic cue, there is no single intuitive process type to which a particular process token belongs, and Conee and Feldman (1998) dispute Heller's (1995) proposal, which has affinities with Olsson's.[4]

So far, the question of whether there are intuitive belief-forming process-types has only received philosophical speculation; the issue has not been settled experimentally. This paper rectifies this by presenting the results of a series of experiments aimed at determining whether most people actually converge on types for normal belief-forming process tokens – and, if so, to what extent reliability-estimates for these types track the corresponding justification and knowledge judgments. Following these experiments is a detailed discussion of several different ways in which these results might be recruited in reliabilist theorizing.

Before the experiments are introduced, the notion of a basic-level concept will be clarified since it is closely connected to the notion of an intuitive process-type.

1 BASIC-LEVEL CONCEPTS

The research into basic-level words and concepts was initiated by Rosch, Mervis, Gray, Johnson, and Boyes-Braem (1976).[5] They investigated taxonomically related concepts: i.e., concepts linked together in a tree in such a manner that the extension of each concept is a subclass of the extension of its parent (if it has one), and the extensions of sister concepts are disjoint. The parent of a certain concept is its *superordinate*; its child is its *subordinate*. For example, in a simple zoological taxonomy containing the concepts 'animal', 'dog', and 'Labrador', 'animal' is the superordinate of 'dog' and 'Labrador' is its subordinate.[6]

It turns out – across several conceptual taxonomies and several semantic domains – that one level of concepts, *the basic level*, has privileged status. For example, basic-level concepts are overwhelmingly preferred in free naming (Rosch et al. 1976; Downing, 1980; Tversky and Hemenway, 1983; Morris and Murphy, 1990); basic-level concepts tend to have significantly more features associated with them than superordinates but not significantly fewer than subordinates (Rosch et al. 1976; Tversky and Hemenway 1983); basic-level categories tend to include things with a common shape (Rosch et al.

4 Here is a quote indicating Feldman and Conee's position.

> Adler and Levin may be suggesting that that belief is judged to be knowledge provided that *the* intuitive process type for the belief is judged to be reliable. This proposal is undermined by the fact that there is no such thing as *the* intuitive type for many beliefs that are judged to be knowledge.
>
> (Feldman and Conee 2002: 104, emphasis in original)

5 Their work was preceded by that of e.g. Brown (1958) and Berlin, Breedlove, and Raven (1973), who dealt with similar issues but used different conceptual and methodological apparatuses.

There is no clear, generally adhered to, distinction in the literature between basic-level *words* and basic-level *concepts*. As is quite common in cognitive psychology, talk of words and concepts are often run together. Whenever possible I will use 'basic-level concepts' in describing the literature.

6 Strictly speaking, in this taxonomy 'Dog' is also subordinate to 'animal' and superordinate to 'Labrador'. However, in the three-level taxonomies often used in the literature, only the top level is referred to as the 'superordinate level' and only the lowest level is referred to as the 'subordinate level'. This paper will follow suit.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

1976); basic-level concepts are the first acquired by children (Anglin, 1977); and basic-level words are more frequent in text (Wisniewski and Murphy, 1989). Of particular interest for present purposes is the research on free naming. It has been demonstrated that people tend to converge on reports at the same level of abstraction when asked to provide names or short descriptions for stimuli, under a variety of circumstances. Convergences have been found in several semantic domains, such as physical objects (Rosch et al. 1976; Downing, 1980), events (Rosch, 1978; Morris and Murphy, 1990), personality traits (John, Hampson and Goldberg 1991) and scenes (Tversky and Hemenway, 1983). In addition, convergences have been discovered using a range of different kinds of stimuli: e.g., photographic pictures, written stories, fake diary excerpts, personal memories of a person or event, and movies. Sometimes the convergence is very striking; Rosch et al. (1976) found that, out of 540 responses, 530–533 converged on the same word for naming a physical object. Sometimes the convergence is slightly less pronounced; but in all of the cited studies, at least 50% of the participants converged on a response at the same level of abstraction.[7]

Although the literature now covers a large semantic domain, it has been almost exclusively concerned with nouns, and nothing in the vicinity of the verbs for belief-forming processes has been investigated. One interesting possible byproduct of the investigation of some of the empirical assumptions in the philosophical debate on the generality problem will thus be a generalization of the notion of a basic level.[8]

2 THE EXPERIMENTS

What follows is a series of three experiments aimed at resolving the conflict between Olsson on the one side and Conee and Feldman on the other. The first experiment is a free naming experiment where participants are asked how people in short film-clips form certain beliefs. The second experiment is a replication of the first, and the third experiment is a rating experiment where one group of participants are asked to judge the degree of reliability of the types they take certain belief-forming process-tokens to fall under, while two other groups are asked to respectively judge the degree of justification of beliefs produced by these process-tokens or whether or not they constitute

7 There are, of course, situations in which people choose not to use the words on which people typically converge. First, a subordinate word might be used to single out an object situated among other objects belonging to the same basic level (Cruse, 1977; Rosch, 1978). Second, a subordinate word might be used when the token is a very atypical member of the higher level (Downing 1980). Third, a subordinate word might be used if the speaker possesses pertinent expertise (Tanaka and Taylor, 1991). Fourth, superordinates might be used instead of lower-level terms when people offer generalizations – so that their statements become as general as is warranted (Cruse 1977). Fifth, a superordinate term might be used when a person is unsure which of several lower-level terms is correct, or when she does not know any term for an object at a lower level (Downing 1980). Sixth, if super- or subordinates have been used at a prior point in the conversation, they might be used again in order to align the dialogue (Pickering and Garrod, 2004). None of these exceptions detract from the way that people in the same situation (i.e., in the same perceptual situation with respect to contrasts, possessing the same level of expertise, etc.) are likelyto converge on the same level.

8 The issue of whether existing explanations of basic-level advantages can be extended to cover belief-forming processes will, for reasons of space, be dealt with separately, in two companion papers. See Jönsson (forthcoming; submitted).

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

knowledge. The experiments aim to establish the validity of the conflicting empirical claims made by Olsson and Conee and Feldman. Once the results have been reported, the exact significance of these for reliabilism will be investigated in detail in the general discussion.

Since complete convergence in verbal reports in any experimental setting is an unlikely result (due to various forms of noise), and neither Olsson nor Conee and Feldman are very precise in their predictions about degree of convergence, something needs to be said about how much convergence that is really needed to validate Olsson's position (and thus conversely, how much divergence that is needed to validate Conee and Feldman's position). One issue that pertains to this is the fact that, from the perspective of the generality problem, what is important is to avoid underspecification. In order to avoid this it is sufficient that participants converge on types with the same reliability, not that they converge on exactly the same types.[9] For as long as participants converge on similar (from the perspective of their reliability) but non-identical types, the reliabilist theory formulated in terms of intuitive types will have clear predictions. This means that if Experiments 1 and 2 reveal that there is sufficient convergence for it to be likely that these types would be assigned the same reliability, then this will be taken as evidence for Olsson's position; if not, this will be taken as evidence against it.

Another issue pertaining to degree of convergence is connected to the literature on basic levels. This literature is almost exclusively concerned with lexical rather than phrasal convergence. So in order to properly compare the results from Experiments 1 and 2 with those from the basic-level literature we should place a special emphasis on the main verb of the participants responses and see to what extent they converge on it. In particular, if at least half or more of the participants converge on the same verb then we have an effect comparable to those found in other domains.

3 EXPERIMENT 1

The aim of the first experiment was to determine the extent to which people converge when they are asked to type belief-forming processes.

3.1 Method

3.1.1 Participants

Fifteen undergraduates from City University in London were paid for their participation.

3.1.2 Stimuli

The process-types that occurred in one or more of a few central texts on the generality problem were first extracted, in order to generate situations to be classified by participants.[10] In order to avoid verbs and verb phrases not in common usage, all expressions

9 Olsson (forthcoming: 2) notes this as well.

10 For this purpose, I made use of Goldman (1979), Feldman (1985), Conee and Feldman (1998), and Feldman and Conee (2002).

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

Table 1. Verbs and verb phrases denoting belief-forming processes

Intuitively Reliable	Intuitively Unreliable
See	Guess
Remember	Dream
Listen	Speculate (V)
Hear	Scan
Feel	Hallucinate
Smell	Have a hunch (V)
Taste	Divine
Reason	Use clairvoyant powers

were excluded that did not have positive frequency counts in the spoken English section of The Corpus of Contemporary American English (COCA). Then, in order to obtain common expressions that might not have been in the original sample, the expressions were run through WordNet, in order to obtain coordinate (sister) terms that also refer to belief-forming processes. All expressions associated with a positive COCA-count for spoken English were retained. A few verbs were assigned new counts, because the initial counts were due to a conflation of counts for several different syntactic categories: e.g., the verb 'divine' was initially assigned a much higher count due to the frequency of the adjectival use of the word.

In anticipation of the third experiment, the resulting set of verbs and verb-phrases was divided into two subsets depending on whether or not the corresponding process was deemed intuitively reliable. The set with the fewest members (the unreliable items) contained eight expressions. These, together with the eight members from the other set with the highest COCA count, can be seen in Table 1.

One of the issues to be tested was whether there is convergence on verbs for belief-forming processes even if no process has been mentioned explicitly (since Conee and Feldman makes a point out of this). Therefore, verbal descriptions of behavior were unsuitable to use as stimuli, even though such descriptions had been used in other places in the basic-level literature (cf. Morris and Murphy's (1990) work on events). Instead, short, high quality film clips were used. This had the benefit of having stimuli that approximated real-life situations - stimuli that featured a range of multimodal properties that could be employed in classification.[11] The television shows *CSI* and *CSI: Miami* were chosen as sources, as it was deemed likely that these shows would contain a number of scenes corresponding to a variety of belief-forming processes. (A recurring theme in both shows is the contrast between proper evidence and more dubious sources of information).[12] Transcripts from the shows were searched for the items in Table 1. Scenes were cut from individual programs corresponding to the first search-hit for each item in Table 1 that fulfilled a number of criteria. These criteria included realism, semantic appropriateness, uniqueness and accessibility of the belief-forming process token, absence of linguistic

11 The stimuli thus have the desireable property of being quite far from the typical short, decontextualized descriptions of possible cases that are common in philosophy.

12 Use of clips from *CSI* and *CSI: Miami* – courtesy of CBS Broadcasting Inc.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

cues,[13] and absence of hard-to-codify objects.[14] The final sixteen clips averaged ten seconds in length.

3.1.3 Procedure

The clips were randomized; half of the participants were shown them in one order, and half were shown them in the reverse order. Four clips were prefixed to the sixteen target clips to avoid any warm-up effects. All participants saw the same four items, though they were not told that they were warm-up items. All participants conducted the experiment on their own in a computer cubicle, where they were left to play clips themselves from within a Powerpoint-presentation.

After having finished typing, participants were asked to fill in reliability-estimates for the processes they had listed. The results of this task are reported separately in Experiment 3.

After completing the experiment, all subjects were debriefed. First, they were asked whether they had had any trouble completing the assignment; and second, whether they were familiar with the TV shows and, if so, whether they thought their familiarity had influenced their answers. No subjects reported having trouble completing the assignment or finding that their familiarity with the shows influenced their answers.

3.1.4 Instructions

The participants were given short booklets to fill in. The first page instructed them to answer questions concerning which process the person in a movie clip underwent in order to arrive at a certain belief. For instance, one clip featured the question 'How did Dominic come to believe that there was a bomb in the school?'. The participants were told to respond as they would have if a friend of theirs that hadn't seen the clip asked them about it.

During piloting, some participants misunderstood the task and answered the question 'why did *x* come to believe that *p*' instead of the question 'how did *x* come to believe that *p*': i.e., they provided facts about the world that did not feature the believers at all. To remind participants of what we were after, each blank to be filled in by the participants was preceded by an appropriate pronoun ('he' or 'she') and three dots.

3.2 Results

Before being analyzed, the participants' responses were formally regimented, to bring out those semantic similarities and differences relevant to the study's objectives. For instance,

13 Scenes were almost invariably discarded if they contained explicit mention of the relevant token-process. That said, in many real-life situations, the best evidence for someone using a certain belief-forming process – e.g., having a hunch or speculating – is that the person *says* she used that process. It would damage the representative character of the sample if items were removed simply because no clips of these tokens could be found that didn't contain verbal mention of them. In the end, fourteen clips were found that had no verbal mention of the relevant token. The two items where this was not possible are marked with (V) in Table 1.

14 An earlier version of this manuscript featured more precise information concerning these criteria but the relevant passages were edited out for space reasons. (The version that contained all the experimental details, was over 17 000 words.) The full details are available on request.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

Table 2. Convergent Responses (verbs)

Item	Conv. Response	Prop.	
Hear	Hear	100%	(15/15)
Have a hunch (V)	Have a hunch	100%	(7/7)
Taste	Taste	87%	(13/15)
Listen	Hear	86%	(12/13)
Smell	Smell	80%	(12/15)
Use clair. powers	Have a vision	73%	(11/15)
Guess	Guess	69%	(9/13)
Feel	Feel	67%	(10/15)
See	See	64%	(9/14)
Scan	See	60%	(9/15)
Remember	Remember	55%	(5/9)
Dream	Dream	54%	(7/13)
Speculate	–	40%	(4/10)
Reason	–	29%	(4/14)
Hallucinate	–	29%	(4/14)
Divine	–	25%	(3/12)

spelling errors and ungrammatical constructions were corrected wherever the intention of the participant was obvious, and whenever an answer was expressed as a conjunction, tacit parts of the conjuncts were, for comparative purposes, spelled out. For example, 'he heard and trusted the information on radio' was changed into 'he heard the information on radio and trusted the information on radio'.[15] A few answers were removed as part of this procedure. These were mainly answers that clearly reflected a misunderstanding of what happened in the clip or what kind of information that was requested. The majority of removed answers were responses to just four items: the *speculate* item, the *have-a-hunch* item, the *reason* item and the *remember* item. These items will be analyzed individually below.

For twelve of the sixteen items, over half the respondents converged on the same verb. Of these twelve items, respondents converged on the expected verb in nine cases and on an unexpected verb in three. Table 2 lists these twelve items, together with the verb or verb phrase that most of the respondents converged on as well as the proportion of respondents who converged on that response. Table 2 also contains, for completeness, the four items where the majority of respondents did not converge on the same verbs.

For both the *scan* item and the *see* item, a few participants used 'looked' (for/at/around) instead of 'see'. These answers were treated as semantically distinct from the *see* response on which most people converged. If these responses are combined with the *see* responses, the proportion of subjects who converged on the same response would have been 78% and 86% respectively (rather than 60% and 64%).

Where respondents used a transitive verb, which they did for eleven of the twelve items, half or more of them converged on the same direct object. For most items, this was a deictic element – such as 'it' – referring to the attributed belief: e.g., how did Eric come to believe that there was a flashlight on the table?...he saw *it*. The only item among the

15 The full details on the regimentation process have been edited out, they are available on request.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

Table 3. Divergent Responses (verbs)

Item	Responses
Hallucinate	See (4) Be on drugs (4) Be high (2) Look at (1) Suffer from delusions (1) Be ill (1) Be shocked (1)
Divine	Believe (3) Know (2) Base on Astrology (2), Be told (2) Remember (1) Learn (1) Have experience with (1)
Speculate (V)	Speculate (4) Suspect (2) Know (2) Make a connection (1) Gather info (1)
Reason	Know (4) Assume (3) Believe (2) Hypothesize (1) Infer (1) Deduce (1) Reason (1) Use his knowledge (1)

twelve where respondents did not converge on precisely the same direct object was the *felt* item, where about half the participants answered 'it' (meaning that there was something in the bag) and half answered 'the bag'.

Overall, the most popular answer was 'he/she V + [PST] it'.

The variability found that was not due to choice of verb or direct object resulted from participants adding adverbials of different kinds (mainly optional prepositional phrases) or conjoining the main clause with clauses describing additional processes. With the exception of responses to the *listen* item, neither adverbials nor extra conjuncts were used by a majority of participants for the items in Table 2. For the *listen* item, the optional prepositional phrases, given by ten out of thirteen subjects,were similar, but signal a significant form of divergence since six of these used 'the radio', while the other four used 'the police radio', and this is a difference that might influence reliability-estimates. The listen-item should thus not be counted as an item where the participants converged on the same process-type.

In order to check whether the order in which items were presented influenced respondents' responses, the answers were split into two groups, one corresponding to each ordering. For most items, both groups exhibited convergence. For the five items where only one group converged on a particular verb, it was investigated whether the items prior to that item for the group that converged contained any mention (either in the clips themselves or in the questions concerning the clips) of the verb, so that the convergence could be explained in terms of a prior verbal cue. No such verbal cues could be found.

For the four items where participants did not converge on the same verb – the *divine* item, the *reason* item, the *hallucinate* item, and the *speculate* item – both verbs and their direct objects were highly varied. See Table 3 for a summary of the verbs used.

Of the three items where participants converged in an unexpected way – the *scan* item, the *listen* item, and the *use clairvoyant powers* item – one item, the *use clairvoyant powers* item, featured a clip where the response that subjects converged on was explicitly mentioned. In this clip, one of the characters is described several times as *having visions*. For the other two items, the response that participants converged on was not explicitly used in the corresponding clips.

3.3 Discussion

Convergence on belief-forming processes was seen for the majority of items with most of the respondents. Respondents converged on the same verbs and, for transitive verbs, the same direct objects; most of them chose not to qualify their reports by adding adverbials

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

or extra clauses describing additional processes. This was true both for presumed reliable and presumed unreliable processes. Since most clips did not have explicit linguistic cues concerning how to type them, Conee and Feldman's (2002) suggestion that usually the only way in which a process type can become salient is through prior mentioning is incorrect. Note that one of the items – the *speculate* item – that did contain an explicit mention of the belief-forming process did not lead to convergence. Therefore, verbal mention is neither necessary nor sufficient to induce convergence.

These findings support the conclusion that a basic level for words for belief-forming processes exists that is similar to what has been found earlier in other semantic domains. Besides the convergence for most of the items, support for this conclusion comes also from the unexpected convergence on the *scan* and *listen* items. If scanning is understood as a form of seeing and listening as a form of hearing, then it seems reasonable to think that the more general items are preferred for structural reasons: i.e., they are understood as having significantly more features than superordinates ('perceive' for example), while subordinates are *not* understood as having significantly more features than them (e.g., listening and hearing are not very distinct kindsof activities). Such preferences have been argued to be typical of basic levels (see e.g. Rosch et al. 1976), lending additional support to the conclusion that there might exist a basic level for belief-forming processes structurally similar to what has been found in other domains.

One item – the *feel* item – contained minor variations in the form of the direct object. Since the semantic differences are so minor, it is very unlikely that they would result in any difference in the perceived reliability of the corresponding processes.

Some items required individual analysis; two items – the *have a hunch* item and the *remember* item – generated many responses that had to be removed, and four items – the *hallucinate* item, the *reason* item, the *speculate* item, and the *divine* item – did not generate convergent responses.[16]

The need to remove answers from the *have a hunch* item and the *remember* item stemmed from the fact that some participants quite clearly answered other questions than the ones they had been explicitly given. For instance, of the participants who provided confused answers for the *have a hunch* item, two were confused about which person in the clip the question was about. Similarly, for the *remember* item, two participants seems to have confused the question 'how did *x* come to believe that *p*' with the question 'how did *you* come to believe that *x* believed that *p*'. These responses, as well as the responses that were confused in other ways, are likely due to idiosyncrasies of the clips rather then with the relevant processes. In particular, the *have a hunch* clip was the only clip that featured a question of conditional structure ('How did Greg come to believe that there *would* be a hit on *y*'s bank trace') and one of very few where both persons in the clip were of the same gender (thus making the name-person association a bit harder). In the *remember* clip, the outward sign (a groan) of the inner sudden realization that signaled remembering, might not have been sufficiently well articulated. As a result, some of the participants might have been at a loss for how to answer the question.

16 The following two paragraphs only contain the main conclusions of the analysis. The full analysis (about 2000 words) was cut for space reasons since it mainly identified shortcomings with the relevant clips and not anything very interesting concerning the relevant processes. The full analysis is available on request.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

Concerning the items where people did not converge in their responses, the main issue seems to be the presence of too many different belief-forming process tokens. The variability thus stemmed from different participants selecting different tokens to type rather than typing the same token differently. Two clips (the *hallucinate* item, and the *speculate* item) also featured information that might have acted as distractors for participants natural responses. In the *hallucinate* clip one of the persons states that the believer (the person whose belief forming process participants are asked about) 'is. . .hyped up on something'. Because the expression 'hyped up' is slang – the exact meaning of which the participants might not have known – they might have hesitated to repeat the phrase verbatim, and instead resorted to paraphrasing. This might mean that the utterance acts as a distractor to participants' natural responses and thus creates divergence. In the *speculate* clip, participants did not report that the believer was speculating even though this clip was one of two clips where the believer herself reported what she was doing. One plausible reason why the participants did not do this is because they didn't believe that she was speculating. In the clip, she is presented as working with the police, and she is sitting in an interrogation room. Participants might have felt that these surroundings warrant a certain degree of certainty on behalf of the interrogating police officer. It also sounds as though she has quite a lot of background information which might have led some participants to use more neutral verbs such as 'conclude'.

The troublesome clips were not possible to replace due to the scarcity of other better clips featuring tokens of the relevant belief-forming processes. However, in anticipation of the replication experiment some minor editing changes to the *hallucinate* clip – the verbal distraction was cut from the clip – and some slight changes to the formulation of the *reason* question – the normal question-format was changed to 'How did x come to believe that p by using information y' to emphasize that the participants should type the process of moving from one information state to another – were made in an attempt to try to circumvent some of the experienced problems.

3.4 *Summary*

Given that the reported patterns are representative of how things are in general, it seems that a reliabilist theory of justification or knowledge can well be constructed that makes reference to the *normal* way in which belief-forming processes are typed. The correctness of this theory depends, in part, on whether the believed reliability of these processes can be used to predict justification and knowledge judgments. This is investigated in Experiment 3. Before that experiment is reported, however, a replication of the first experiment will be described, even though it took place after Experiment 3. The replication experiment is reported first to make the comparison between it and Experiment 1 as clear as possible.

4 EXPERIMENT 2

The primary aim of the experiment was to replicate Experiment 1 and test whether minor modifications to the *reason* item and the *hallucinate* item would result in greater convergence.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

4.1 *Method*

4.1.1 Participants

Eighteen undergraduate philosophy students from Lund University volunteered for the experiment.

4.1.2 Stimuli

The movie clips from Experiment 1 were re-used. Two new items – *read* and *tell* – were added, as these were felt to be commonly used belief-forming procedures, even though the method used to generate stimuli in Experiment 1 did not produce them. The *hallucinate* item and the *reason* item were changed slightly (see the discussion of Experiment 1).

4.1.3 Procedure

Participants (nine at a time) were shown the clips collectively in a classroom, to make the testing more efficient. All clips were played at least two times. Participants were instructed to raise their arm if they wanted to re-play a certain item again. Participants were each given a booklet as in Experiment 1, which they filled out without any interaction with any of the other participants. Clips were played in a new, randomized order.

4.1.4 Instructions

The instructions were the same as those in Experiment 1.

4.2 *Results*

As in Experiment 1, participants' responses were formally regimented to bring out semantic similarities and differences. A few answers were removed. Again, these were mainly answers that clearly resulted either from misunderstanding what happened in the clip or what kind of information was requested. The majority of the removed answers were responses to just two items: the *had a hunch* item and the *reason* item.

Over half the respondents converged on the same verb for fifteen of the eighteen items. Of these fifteen, respondents converged on the expected verb (with respect to the updated expectations that followed from the results of Experiment 1) in all but one case: with the *reason* item, participants converged on 'believe'. Table 4 lists the items on which participants converged, together with the verb or verb phrase that most of the participants converged on, along with the proportion of participants who converged on that response. Also included in Table 4 are the three items where participants did not converge on the same verb.

Like in Experiment 1, the *scan* item and the *see* item received a few answers ('observed' or 'recognize' for the *scan* item, 'observed' for the *see* item) that were very close, but not identical, to the ones on which most respondents converged. If these had been combined with the responses on which most participants converged, the proportions would have been 94% and 78% (rather than 50% and 67%).

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

Table 4. Convergent Responses (verbs) in Experiment 2

Item	Response	Prop.	
Hear	Hear	94%	(17/18)
Taste	Taste	83%	(15/18)
Listen	Hear	83%	(15/18)
Read	Read	78%	(14/18)
Smell	Smell	72%	(13/18)
Use Clair. powers	Have a vision	72%	(13/18)
Feel	Feel	72%	(13/18)
Reason	Believe	71%	(10/14)
See	See	67%	(12/18)
Tell	Tell	67%	(12/18)
Have a hunch	Have a hunch	64%	(7/11)
Dream	Dream	53%	(9/17)
Scan	See	50%	(9/18)
Remember	Remember	50%	(9/18)
Guess	Guess	50%	(8/16)
Divine	–	35%	(6/17)
Hallucinate	–	33%	(6/18)
Speculate	–	35%	(4/17)

Where the respondents used a transitive verb with thirteen of the fifteen items, over half converged on the same direct object. For many items, this was just a deictic element – such as 'it' referring to the content of the attributed belief. The two items where respondents did not converge on precisely the same direct object were the *hear* item and the *reason* item. For the *hear* item, the direct objects were still largely convergent: most respondents used either 'a woman's scream' or 'a scream'. For the *reason* item, where respondents converged on 'believe', the direct objects were highly varied. All included details of what the character had said explicitly, but they differed both in how much information was included and which parts of the available information was included.

Overall, the most popular answer was 'he/she V + [PST] it'.

The variability found that was not due to choice of verb or direct object resulted from participants adding adverbials of different kinds (mainly optional prepositional phrases) or conjoining the main clause with clauses describing additional processes. With the exception of responses to the *listen* item and the *read* item, neither adverbials nor extra conjuncts were used by the majority of participants for the items in Table 2. For the *listen* item, the prepositional phrases used were very similar to those found in Experiment 1; some respondents used 'radio' and others used 'police radio'. This should, like in Experiment 1 be taken as a significant form of divergence. The same goes for the *read* item where the prepositional phrases reported were quite divergent, including 'on a paper', 'off the autopsy report', 'on a journal about the victims', 'in the report', 'in the coroner's record', 'from the protocol', and 'in a report about the dead body'.

For two of the items where participants did not converge on the same verb – the *divine* item and the *speculate* item – both verbs and their direct objects were again quite varied. However, responses for the *reason* item and the *hallucinate* item were more convergent with respect to the verbs used than in Experiment 1. In the hallucinate clip, four

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

participants now reported that the man in the clip hallucinated, and nine persons either said that he saw or that he thought he saw the man's face melting.

4.3 Discussion

The results from Experiment 1 were replicated: there was convergence on belief-forming processes for the majority of items. Respondents mostly converged on the same verbs and, with transitive verbs, on the same direct objects. Most of them chose not to qualify their reports by adding adverbials or extra clauses describing additional processes. The addition of two new items – *read* and *tell*, both of which elicited convergent responses – meant that participants now converged on a slightly larger proportion of the items tested: participants converged on the same verbs on 83% rather than 75% of the items. Three items (the 'listen'-item, the 'read'-item and the 'reason'-item) differed either in their direct objects or in their optional prepositional phrases. Given that these items were counted as significantly divergent, the general convergence on process-types would be 67% and 68% respectively.

Two out of the three items where participants did *not* converge on the same verb – the *speculate* item and the *divine* item – retained the problems identified in the previous discussion: too many tokens and conflicting linguistic evidence. These problems are probably responsible once again for the divergences found. The third item – the *hallucinate* item – now showed greater convergence. If one assumes that those respondents who answered that the man in the *hallucinate* clip 'saw' something meant this in a non-veridical sense, then exactly half the respondents now converged on the same verb. The hypothesis that the verbal cue included in the earlier version of the clip acted as a distractor to respondents' natural response is thus given some support.

The *reason* item was also changed to try to circumvent what was assumed to be a problem with too many tokens of belief-forming processes in the clip. The question was changed so that it should elicit a response in terms of the kind of mental transition – e.g., reason – that took place. Although the answers were now convergent, they did not converge in the expected direction, and the direct objects used were highly varied. It is possible that the variability can be accounted for thus: because the direct objects of propositional-attitude verbs are propositions, the classification task subjects engaged in – given that they wanted to attribute a belief – amounts, to a large extent, to selecting what of everything the believer said should be included in the attribution. If this is right, the situation is probably best seen as one where there really are multiple tokens present, since there are a number of different beliefs.

The variability found among the prepositional phrases used for the *read* item likely arises because participants found the clipboard from which the character reads quite hard to describe – especially since they were non-native English speakers. This form of variability is thus akin to that found by Downing (1980) for hard-to-codify objects.

To sum up: it seems likely that the few divergencies that did show up in this and the previous experiment can be traced to a few, already known, spurious causes of divergent naming: in particular, the presence of objects that are hard to codify, the presence of several different process tokens, and the presence of distracting verbal cues that participants do not feel compelled to repeat either due to use of slang or the presence of contravening evidence. In the absence of such factors people converge on the same types for belief forming processes.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

5 EXPERIMENT 3

Even though people seem to classify belief-forming processes in the same way, the generality problem for reliabilism is not thereby solved. It also needs to be the case that the relevant process-types are reliable – at least roughly – in those circumstances (and only those circumstances) where persons judge beliefs to be justified, and those circumstances where persons judge beliefs to be knowledge. The aim of the third experiment was to investigate this.[17] The experiment was carried out in conjunction with Experiment 1.

5.1 Method

5.1.1 Participants

Forty-five undergraduates from City University in London were paid for their participation. Fifteen of these were the participants from Experiment 1.

5.1.2 Procedure

Participants were randomly divided into one of three conditions: a reliability-condition, a justification-condition, and a knowledge-condition. Participants in each group were asked to give judgments concerning either the reliability of belief-forming processes (the reliability-condition), the extent to which people were justified in believing something (the justification-condition), or whether or not someone knew something (the knowledge-condition). Participants in the reliability-condition completed a typing-assignment (the one reported in Experiment 1) before being asked to provide reliability-judgments. Participants in the justification and knowledge-conditions started out directly by giving justification or knowledge-judgments, since if the participants in the justification and knowledge-conditions had also, first, given reliability-estimates or typed processes, there would have been the risk that participants had conflated justification-ratings, or some other epistemic property of the identified processes with the one they were explicitly asked for. So in order to avoid this confounding factor, correlations were tested across participants.

Participants in the knowledge-condition completed an additional task after the knowledge-judgment task. The motivation is as follows: knowledge-reliabilism, like most theories of knowledge, entails that the believed proposition is true, and one therefore needs to make sure that the reason a participant judges a belief not to be knowledge does not have to do with the participant believing the relevant proposition to be

17 Doing this in an experimental setting will hopefully alleviate a worry that some critics of reliabilism have expressed. For example, Feldman (1985) seems to suggest that reliabilism sometimes appears more attractive than it really is because, when going through examples, the reliabilist might have a tendency to type so that reliabilism delivers the intuitively correct results – even though she has no principled grounds for her typing. When faced with one justified belief and one unjustified belief, the reliabilist might tacitly type so that the first results from a reliable process and the second from an unreliable process. This will strengthen the reliabilist's conviction in the truth of reliabilism, even though it might be the case that no principled way allows one to type so that both cases turn out as desired. Eliciting typing and reliability judgments sequentially in the way described here avoids this kind of bias, as does the between-subjects design.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

false.[18] The additional task was included to determine this. It simply involved the participants judging whether the attributed propositions in the clips were true or false, and if they were false, whether their being true would have been sufficient (in the circumstances in the clips) for the participant to feel that the person in the clip had knowledge.

After completing the experiment, all subjects were debriefed by asking them, first, whether they had any trouble completing the assignment; and, second, whether they were familiar with the TV-shows and, if so, whether they thought that their familiarity had influenced their answers. No subjects reported that they had had trouble completing the assignment or that they thought that their familiarity with the shows had influenced their answers.

5.1.3 Instructions

Participants in the three conditions were given separate instructions. The participants in the reliability-condition were asked to indicate on a scale from one to seven 'how reliably you think the processes, in circumstances similar to those in the clip, generate true beliefs'. The participants in the justification-condition were asked to indicate on a scale from one to seven 'how *justified* you think the person in the clip is in believing a certain belief, that is, how good reason that person has to believe that belief.' Finally, the participants in the knowledge-condition were asked whether they 'think it is appropriate to say of the person in the clip that he or she *knows* a certain proposition.'

5.1.4 Stimuli

The movie-clips from Experiment 1 were re-used, but the subjects in the justification and knowledge-conditions received different questions.

5.2 *Results*

One participant in the knowledge-condition failed to follow the instructions and was excluded.

The means for the individual items and the across-item mean were calculated for the reliability and justification conditions. The across-item mean for the reliability judgments was 4.77 ($SD = 0.98$), and the across-item mean for the justification judgments was 4.57 ($SD = 1.19$). The means for the individual items are presented in Table 5. A Pearson's r test revealed a very strong correlation between reliability and justification-judgments: $r(14) = .84$, $p < .0005$.

For the knowledge-condition, items were classified as knowledge^{+} if all participants, or all but one, agreed either that they were knowledge or that they would be if the relevant proposition was true. This was true of the *see* item, the *listen* item, the *taste* item, the *smell* item, the *reason* item, the *scan* item, the *feel* item, and the *remember* item. A threshold-value between 4.54 and 5.19 on the mean reliability correctly predicts fourteen out of sixteen items. It incorrectly classifies the *hearing* scenario as one of knowing and the

18 Many reliabilist positions include additional requirements as well, and such requirements could at least theoretically also constitute confounds. But since this seemed less likely, no measures corresponding to that taken for the truth requirement were taken to remove these possible confounds.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

Table 5. Reliability and Justification Judgments

Item	Mean Rel.	Mean Just.
See	6.13 (1.06)	6.40 (1.54)
Listen (Hear)	5.80 (1.21)	6.20 (1.31)
Remember	5.73 (1.28)	4.87 (3.70)
Taste	5.67 (1.50)	5.64 (3.37)
Feel	5.60 (1.59)	4.87 (1.98)
Smell	5.47 (1.41)	5.33 (4.95)
Hear	5.47 (1.46)	5.40 (4.54)
Scan (See)	5.20 (2.14)	4.87 (4.12)
Speculate (NA)	4.53 (1.68)	4.60 (2.83)
Hallucinate (NA)	4.33 (2.29)	2.67 (3.10)
Reason (NA)	4.27 (1.22)	5.07 (3.50)
Have a hunch	4.20 (1.82)	4.27 (3.50)
Dream	3.87 (2.29)	3.73 (3.64)
Guess	3.60 (1.80)	2.33 (3.38)
Divine (NA)	3.20 (2.37)	3.13 (3.12)
Use Clairvoyant powers (Have a vision)	3.20 (2.18)	3.73 (4.07)

reasoning scenario as one of not knowing. An optimal threshold-value on the justification-means – a value between 4.54 and 5.72 – also got *hearing* wrong but got the rest of the items right.

Since the unreliable processes in Experiment 1 generated a slightly lower mean degree of convergence than did the reliable processes, it is important to rule out that the found correlations are not just effects of reliable and unreliable processes generating different degrees of convergence.[19] A Spearman's *r* reveal that the correlation between the ranked order of degree of convergence and ranked degree of justification is actually quite weak $r_s = 0.44$, which is not significant as a one-tailed test at $p = 0.01$ (and barely significant at 0.05), a result very distant from the found reliability-justification correlation. It can thus be concluded that the correlation described above was not due primarily to the slight differences in the degree of convergence between reliable and unreliable types found in Experiment 1.

In order to verify to what extent the participants reliability-estimates clustered together (and thus to what extent it was likely that participants assigned reliability to the same types, see below for the relevance of this) a split-half reliability-testwas performed by randomly dividing the reliability-estimates for each item into two sets and then measuring the degree of correlation across items between these sets. A very strong correlation was found $r_{sb}\ (14) = .87$, $p < .0005$.

5.3 Discussion

The two main results from this experiment support the claim that reliability-estimates assigned to belief-forming process-types can do a very good job of tracking justification and knowledge-judgments. Moreover, the reliability-estimates assigned by the participants

19 This was pointed out by an anonymous reviewer for this journal.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

to each item were fairly similar across participants. The exact significance of this will be discussed below.

6 GENERAL DISCUSSION: RELIABILISMS BUILT ON CONVERGENCE

The reported experiments lend quite strong support to Olsson's (forthcoming) position on intuitive process-types; contrary to the claims made by Feldman and Conee (2002), people often converge on the same types when reporting process-tokens even in the absence of explicit linguistic cues, and it thus makes sense to speak of *the* intuitive process type of a process token. Expressions for belief-forming procesess thus exhibit the kind of convergence that is characteristic of basic-level words.[20] In addition, the reliability estimates reported in Experiment 3 could be used to predict justification and knowledge judgments.

Although these results might well underwrite a solution to the generality problem, Olsson (forthcoming) is somewhat unclear about the the details of the kind of proposal he had in mind, and it is not obvious how the reliabilist makes optimal theoretic use of these results. In particular, it is not clear which types that Olsson considers to be the ones whose reliability is important for justification and knowledge. When turning to the issue of convergent typing he first describes that he anticipates (pending something like the present investigation) that people will '*report* in a homogeneous manner' (p. 10, emphasis added), and thus suggests that the relevant types are those that are explicitly expressed by people. Later on in his discussion, however, he expresses his prediction in terms of the types that people will 'think' or 'classify' in terms of (p. 16), thereby suggesting that the relevant typing is cognitive rather than verbal. So it is unclear if Olsson thinks that the process-types that are relevant from the perspective of justification and knowledge judgments are i) linguistic, or, ii) cognitive. Moreover, it is not clear *whose* types Olsson thinks are important for justification and knowledge. In particular, is it i) those that people in general tend to converge on (i.e. the intuitive types), ii) those of the believer (i.e. the person whose belief is being evaluated in terms of justification or knowledge), or iii) those of the attributor of justification and knowledge. The first interpretation fits the discussion best but presupposes that a certain interpersonal uniformity must obtain in order for someone to be justified or have knowledge, an unintuitive position. The latter two positions doesn't have this problem but have a less immediate connection to the convergence results.

Altogether, six different possible reliabilist theories can be generated from the aforementioned distinctions.[21] Each of them will be scrutinized in this section. The theories

20 Whether or not these words also exhibit other properties associated with basic-level words, or whether or not the found convergence can be explained in the same way as that of other basic-level words, is investigated in two companion papers. See Jönsson (forthcoming; submitted).

21 There are additional distinctions that can be drawn as well, but these have less impact in the present context. For instance, we will see below that one of the theories can be plausibly framed both in terms of *actual typing*, and in terms of *dispositional typing*. Moreover, all theories can be framed both in terms of *estimated reliability* and *actual reliability* (cf. Goldman1979). Concerning the latter distinction, my default below is estimated reliability since this is closest to the data. However, one could also frame the theories in terms of actual reliability without severing the ties between the theories and the results of Experiment 3. The way to do this is just to assume that participants estimates of reliability were more or less correct. But without any independent verification of this, this option seems much less plausible than the alternative.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

will both be evaluated from the perspective of how well the data reported in this paper supports their empirical presuppositions, and from the more philosophical perspective of how plausible they are as theories of justification and knowledge.[22] Each subsection below will start with a precise statement of a theory, followed by a description of the support, if any, its empirical presuppositions get from the reported results, and a discussion of possible problems with the theory.

Proposal 1: Use the Intuitive Linguistic Types

(J^{IL}) A believer B is justified in believing p to the degree j iff $j = f(r_{IL}(t_{IL}(bfp)))$,

where bfp is the token-process that led B to believe that p, t_{IL} (bfp) is the intuitive linguistic type to which bfp belongs, i.e. the type in terms of which people in general are disposed to report bfp, r_{IL} is thedegree of reliability that is generally attributed to t_{IL} (bfp), and f is a monotonic increasing function of degree of reliability.[23]

The central empirical presupposition of this theory – that people tend to converge on the same linguistic types for belief-forming processes – is directly supported by the results of Experiments 1 and 2. At first glance it also seems that Experiment 3 demonstrated that the estimated reliability of these processes could be used to predict justification and knowledge judgments. However, there are several problems with J^{IL}.

First, the proponent of J^{IL} is forced to say that if people in general do not converge in their dispositions on how to report some particular process token, then J^{IL} doesn't make any predictions about justification, on pain of its false presupposition. In fact, since J^{IL} has the presupposition it has, it follows that if people actually did not converge on reports for belief-forming processes in general, then talk of justification would be meaningless. This peculiar dependence of J^{IL}'s makes it very implausible (independently of whether, as a matter of fact, the theory has any truth-value gaps). The fact that J^{IL} presupposes a socio-linguistic fact which is largely independent of the particular facts concerning any

22 In order to minimize redundancy I will only discuss theories of *justification* in the main text. For each of these, I imagine the corresponding theory of knowledge to be like the theory of justification but with an added threshold on reliability and a truth condition.

To illustrate, corresponding to J^{IL} below is K^{IL}.

(K^{IL}) A believer B knows that p iff: i) B believes that p, ii) p is true, and iii) $r_{IL}(t_{IL}(bfp)) > k_A$,

where bfp is the token-process that led B to believe that p, t_{IL} (bfp) is the intuitive linguistic type to which bfp belongs, i.e. the type in terms of which people in general are disposed to report bfp, r_{IL} is thedegree of reliability that is generally attributed to t_{IL} (bfp), and k_A is some threshold on reliability that A thinks is required for knowledge.

23 The third experiment demonstrates that both justification and knowledge judgments can be predicted-from r_A. However, since reliability and justification judgments were not identical, one must incorporate a function f from r to justification judgments into one's analysis of justification. The exact regression function that can be extracted from the data in Experiment 3 is scale dependent and thus should not be identified with f. Instead, f must be a generalization of that function. Similarly, the exact threshold used to predict knowledge$^+$ judgments from reliability assignments is also scale dependent and should not be used in a general reliabilist theory; instead, a generalization k_A should be used.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

particular instance of justification attribution, needs to be backed up by explanation or argument.

Second, it turned out that that the linguistic types on which people actually converged were quite general (e.g. 'He saw it'). This result poses at least a *prima facie* problem for J^{IL} which was anticipated by Feldman (1985). According to Feldman, general types such as those reported in Experiments 1 and 2 open up for obvious counterexamples to reliabilist theories. For example, even if seeing a distant object and seeing a nearby object are both cases of seeing, a belief about there being, say, a mountain goat at a certain place can be justified if the goat is nearby but might not be if the goat is far away. To take another example, someone can be justified in believing, on the basis of seeing, that there is an animal at a certain place, even if she is not justified in believing, under the same observational circumstances, that there is a mountain goat at that place.

However, this argument makes the following assumption (which we might call *the convergent typing of conspecifics assumption*): if one process-token is typed in terms of a certain type, then every other token of that type must also be typed in terms of this type. For instance, if a certain token is typed in terms of seeing, then every other token of seeing must also be typed in terms of seeing. This principle, or something like it, is clearly at work in both of the two examples above. However, the convergent typing of conspecifics doesn't follow from J^{IL}. Even if people are disposed to report a particular process-token in terms of a certain type, it does not follow that they must also be disposed to report every other token of that type in terms of that type. It is actually a questionable empirical assumption that they do. Independently of J^{IL}, it is unclear why any proposed solution to the generality problem should accord with this principle.[24]

Third, the reliability estimates given by participants in Experiments 3 (which turned out to predict justification and knowledge judgments) are probably best understood as estimates of more narrow types than those explicitly reported, such as the types in terms of which the participants thought about the processes. Participants were encouraged to make estimates of the reported processes in circumstances like those in the relevant clips and are thus likely not estimates of the reliability of these types in the abstract. This can also be seen in the results from Experiments 3 where, for instance, both the scan-item and the see-item were reported in terms of 'He saw it' but they were given different reliability estimates. On pain of inconsistency, this suggests that these estimates must have been estimates of something more fine-grained than the explicit types, such as the types in terms of which the participants thought about the relevant

24 The associate editor assigned to this paper suggested a 'speculative hypothesis' about the general shape of linguistic process-typing which is relevant in this context. According to him/her it might be the case that participants type processes broadly ('he sees it', etc.) when circumstances are typical, but when they're not (e.g., lighting conditions are bad, the object was moving quickly), they type them more narrowly, taking such conditions into account. The associate editor points out that this is compatible with Wegner and Vallacher's Action Identification Theory (2011) where common actions are conceptualized in fairly abstract terms (e.g., eating Cheetos, ringing the doorbell) but in more molecular terms when the actions are unusual (e.g., eating Cheetos with chopsticks, moving my fingerforward)– especially when they're more difficult. Although this is an interesting proposal, whether or not it is true needs further investigation (which is also what the the associate editor suggests). But it can be noted that if the hypothesis holds up it is only relevant for reliabilism if the the convergent typing of conspecifics thesis is off the table. It should also be noted that J^{IL} might not make the best use of this data due to its other problems. See J^{AL} for a stronger thesis.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

processes.[25] This suggests that the reliabilist who is interested in harnessing the predictive power of the reliability estimates from Experiment 3 should not – like J^{IL} does – recruit the linguistic types in her theory but their cognitive counterparts. This is done in the next theory.

Proposal 2: Use the Intuitive Cognitive Types

(J^{IC}) A believer B is justified in believing p to the degree j iff $j = f(r_{IC}(t_{IC}(bfp)))$,

where bfp is the token-process that led B to believe that p, t_{IC} is the intuitive cognitive type to which bfp belongs, i.e. the type in terms of which people in general are disposed to think about bfp, r_{IC} is the degree of reliability that is generally attributed to t_{IC} (bfp).

The central empirical presupposition of J^{IC} – that people tend to converge on the same cognitive types for belief-forming processes – is supported in two ways by the data from the three experiments.

First, the found linguistic convergence from Experiments 1 and 2 supports the existence of cognitive convergence on the assumption that barring evidence to the contrary, convergence in thought is the best explanation of linguistic convergence (i.e. if we describe things in the same way it is likely that we do so because we think about them in the same way). Call this the *CEC*-assumption (the convergence-explains-convergence-assumption) for short. The *CEC*-assumption is of course especially compelling to make in cases where what is said and what is thought is the same,[26] but it also seems to be compelling to make in situations where this is not the case. For instance, consider reports featuring universal quantification. If two people at a party and in otherwise similar contexts both reports that 'Everyone is having fun' it is, barring evidence to the contrary, likely that they have the same thing in mind, even though this cannot reasonably be what they explicitly say (what they have in mind is more likely something like 'Everyone *at the party* is having fun.'). Or take color attribution; if two people in similar conditions, both say, 'Lisa's car is red' it is likely that they have the same thing in mind, even though this must be something like 'Lisa car appears red in lighting conditions like these'. Or take assertions of subjunctive conditionals; if two people in similar situations both say 'If kangaroos didn't have tails they would topple over' it is likely that they have the same thing in mind, even though this is likely to involve, among other things, a certain tacit restriction to a subset of all possible situations. It seems that belief-reporting is like these other forms of reporting and that the *CEC*-assumption is thus reasonable to make in this context. The

25 What is explicitly said, and what is meant or thought, often comes apart in various kinds of reports since speakers normally leave implicit, for instance, information which is contextually salient or difficult to describe. Consider reports featuring universal quantification for instance. At the party, we can drop 'at the party', when we report that 'Everyone is having fun' even though this creates divergence between what we think and what we literally say. Or take color attribution; we say that 'Lisa's car is red' even though the car only appears redin certain lightning condition, but we seldom make explicit this qualification. Or take assertions of subjunctive conditionals; we say that 'If kangaroos didn't have tails they would topple over' even though this is only true given that we consider a certain subset of all possible situations.

26 If we have perfect convergence this assumption is then just the transitivity of identity.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

found linguistic convergence in Experiments 1 and 2 thus give support to the central empirical presupposition of J^{IC}.[27]

Second, the split-half reliability test from Experiment 3 suggests that the participants reliability estimates were clustered fairly tightly together. The simplest explanation of this is that participants are estimating reliability of the same types, since the degree of similarity between speakers' reliability-estimates in a particular situation is an indication of whether or not they have the same type in mind, since people who have the same type in mind are more likely to give similar reliability-estimates than participants having different types in mind. Given that the types that are assigned reliability estimates are cognitive (for the reasons given on the previous page), the split-half reliability test thus supports the central empirical presupposition of J^{IC}.

If we instead turn to the matter of how well the data supports J^{IC}'s predictions of justification and knowledge judgments, Experiment 3 demonstrates that the reliability of cognitive types can be used to predict justification and knowledge judgments very well.

All in all, even if J^{IC}'s central empirical presupposition is supported less directly than that of J^{IL} (and depends on the *CEC*-assumption), J^{IC} still seems to be an empirically plausible theory. However, this theory also face problems, the first of which is particularly damaging.

First, J^{IC}, like J^{IL}, depends on a particular social fact to obtain – that people in general are disposed to think alike about process tokens – in order to make predictions about justification, even though this fact is largely independent of the facts concerning any particular instance of justification attribution. This is very unintuitive.

Second, J^{IC}, unlike J^{IL}, has the *prima facie* problem of not making the relevant process-types explicit. In the literature on the generality problem it is sometimes tacitly assumed that in order to solve the generality problem, the reliabilist has to make explicit the types which she thinks are relevant for justification and knowledge in concrete cases. Call this the *full explicitness requirement*. J^{IC} might be held to violate this requirements since it only tells us in a concretesituation that the type which is relevant is *that which people in general will think about the relevant belief-forming process in terms of*, not if this process is e.g. *seeing*, *seeing in poor lightning conditions*, or *seeing a flashlight a distance*.

However, the full explicitness requirement does not seem to be a reasonable requirement to make in general. The requirement is likely to be motivated by a testability

27 Another way to argue for the *CEC*-assumption is to remark that it is a special case of an explanation of a *kind* of event. In order to explain a kind of event we need to provide information which is common to the causal histories of all tokens of that kind (cf. Lewis 1986). In the present case, we are thus required, in order to explain linguistic convergence, to provide common information from the individual histories of all the convergent utterances (such as the information that the speakers had the same process-type in mind).

Another way to frame this is to say that if two people converge linguistically but have different things in mind, then there is properly speaking no explanation of why they converged. There is one explanation of what led one person to her utterance, and one of what led the other to her utterance, but no explanation of why they converged. So the cases of convergence reported in this paper could either be explained by people having the same type in mind (and having some common reason for not making this kind completely explicit), or be treated as accidents where each individual typing has to be explained individually. The former kind of explanation is, in the absence of additional information, clearly better than the latter, and via an inference to the best explanation, what we should believe.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

consideration; it should be possible for someone other than the proponent of a particular reliabilist theory to test that theory's predictions, and one way to make sure of this is to require that the types are made explicit. However, since the types presupposed by J^{IC} are – by definition – shared, it is likely that the appropriate type to test the theory with in a certain situation by a certain person just is the type in terms of which that person is likely to think about the process. So there is then no need for J^{IC} to also satisfy the full explicitness requirement. If we are likely to converge on the same types in thought, it is unnecessary to tell each other which they are.

It should also be noted that it is not as if J^{IC} tells us nothing about the relevant process-types. For even if the verbal reports of process-types are likely to be incomplete renditions of their cognitive counterparts, the two types are likely to be congruent in the sense of the cognitive types being sub-types of the linguistic ones. For instance, if we report something as a case of seeing, the corresponding cognitive type is likely to be a particular subspecies of seeing. So, J^{IC} tells us something explicitly about which types that are relevant.

Since both J^{IL} and J^{IC} shared the problem of having a peculiar social dependence I will now consider four different theories where this is not the case. The first two can be quite briefly dismissed (but are included in order to provide good contrasts with the final two proposals), but the latter two (in particular the final one) are more promising.

Proposals 3 & 4: Use the Believer's Types

(J^B) A believer B is justified in believing p to the degree j iff $j = f(r_B(t_B(bfp)))$,

where *bfp* is the token-process that led B to believe that p, t_B is the type in terms of which B is either disposed to report *bfp*, or is disposed to think about *bfp*, t_B is the degree of reliability that B estimates t_B (*bfp*) to have. There is no need to differentiate here between linguistic and cognitive versions of this theory since both theories face the same objection; it is not true in general that a believer is able to report her belief-forming process or have any thoughts about it. Small children and animals can have justified beliefs but might not have the resources to talk or think about belief-forming processes.

Proposal 5: Use the Attributor's Linguistic Types

(J^{AL}) A believer B is justified in believing p relative an attributor A to the degree j iff $j = f(r_{AL}(t_{AL}(bfp)))$,

where *bfp* is the token-process that led B to believe that p, t_{AL} is the type in terms of which A is disposed to report *bfp*, and r_{AL} is the degree of reliability that A estimates t_{AL} (*bfp*) to have.[28]

28 This is a species of attributor epistemology, i.e. a form of contextualism that maintains that facts about the attributor of justification and knowledge (such as various psychological facts) are relevant over and above the facts about the subject of attribution (cf. Rysiew 2011).

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

The central empiricial presupposition of this theory – that there is such a thing as the process-type attributors are disposed to report process-token in terms of – is indirectly supported by the results of Experiments 1 and 2, via the inference from interpersonal convergence to intrapersonal stability: if it is likely that the members of a group of people do the same thing x in a certain situation, then the members of that group must each individually be fairly likely to do x in that situation.

J^{AL} avoids the peculiar dependence problem of J^{IL} and J^{IC}, by not invoking interpersonal convergence explicitly (instead, J^{AL} is related to the convergence through an evidential relation). Granted, in situations where the attributor is not disposed to report a process-token in any particular way, J^{AL} doesn't make any predictions about justification. But in these cases the would be attributor might also have trouble judging whether justification or knowledge obtains. Whether or not this is so needs to be tested independently, but this form of dependency on facts about the attributor is, at the very least, understandable, which the dependency on independent social factors is not.

The main problem with J^{AL} is the same as that with J^{AC}; it is likely that the reliability estimates from Experiment 3 belong to cognitive, rather than linguistic, types. The final two theories avoids this problem.

Proposals 6a & 6b: Use the Attributor's Cognitive Types

(J^{AC}) A believer B is justified in believing p relative an attributor A to the degree j iff $j = f(r_{AC}(t_{AC}(bfp)))$,

where bfp is the token-process that led B to believe that p, t_{AC} is either the type in terms of which A is actually thinking about bfp at the time of attribution (J^{ACa}), or the type in terms of which A is disposed to thinkabout bfp at the time of attribution (J^{ACb}), and r_{AC} is the degree of reliability that A estimates t_{AC} (bfp) to have.

J^{AC} comes in two varieties depending on if the relevant cognitive type is the one the attributor is thinking about the process-token in terms of as part of the attribution of justification or knowledge, or that type which an attributor is *disposed* to think about that token at the time of attribution.[29]

The central empirical presupposition of J^{ACb} is supported by the results from Experiments 1–3, by joining together the support for J^{IC} with that for J^{AL}, i.e. Experiments 1 and 2 directly supports interpersonal linguistic convergence, and from this we can infer, given the *CEC*-assumption and the split-half reliability test in Experiment 3, interpersonal cognitive convergence, and from this we can infer intrapersonal cognitive stability along the lines described for J^{AL}.

The central empirical presupposition of J^{ACa} is only supported by the evidence reported here on the assumption that the attributor of justification and knowledge is thinking of the process-token at the time of attribution. If this is the case, then intrapersonal stability of cognitive types makes the central presupposition of J^{ACa} likely. However, the assumption that an attributor is thinking about a process-token as part of attributing justification is

29 Even if the actualist (rather than dispositionalist) proposal doesn't have any particular advantages over its dispositional counterpart, I discuss it here nonetheless for completeness. Note also that there are no plausible actualist counterparts to the proposals that have been discussed above.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

not supported by the reported data, and I'm unaware of any other data that makes this likely. Of the the two variants, J^{ACb} thus has a better supported empirical presupposition, and in the rest of this section I will be exclusively concerned with it.

J^{ACb} avoids all the problems that we have so far discussed. First, J^{ACb}, like J^{AL}, avoids the peculiar dependence problem by not explicitly invoking interpersonal convergence. It might still be subject to presupposition failures, but if it is, there might be a corresponding suspension of justification judgment. Second, since J^{ACb} in a clear sense corresponds to J^{IC}, the full explicitness requirement is not appropriately brought to bear on it for the same reason this was inappropriate for J^{IC}; J^{ACb} is testable even if it doesn't make the corresponding types explicit since the types it invokes are shared among people. Third, since the attributor can be distinct from the subject of attribution (the believer) it doesn't follow from it, like it did from J^{B}, that organisms that cannot think about processes cannot be justified. It does follow from J^{ACb} that one cannot *attribute* justification (in any substantial sense) unless one is able to think about belief-forming processes. Although this prediction also remains to be tested, it is considerably more plausible than the thesis that one must be able to think about belief-forming processes in order to have justified beliefs.

In light of the above considerations I submit that J^{ACb} is the best way to solve the generality problem by recourse to the experimental data reported in this paper.

7 FINAL REMARKS

To sum up the discussion so far; the proposals in terms of linguistic types (i.e. J^{IL} and J^{AL}) are directly supported by the results from Experiments 1 and 2 but might involve only coarser-grained variants of the types whose estimated reliability were shown to predict justification and knowledge judgments in Experiment 3.

The proposals formulated explicitly in terms of a social convergence (i.e. J^{IL} and J^{IC}) rests on a peculiar form of dependence, which even if it doesn't *de facto* threaten the theories' generality, seems unbelievable. J^{ACb} avoids these problems. However, by doing so, it loses the direct empirical support of J^{IL} and J^{AL}. Nonetheless, its empirical presupposition does receive empirical support, albeit indirectly, from the experiments reported here and J^{ACb} thus seems to be a promising way for the reliabilist to come to terms with the generality problem.

There are a number of ways in which empirical presupposition of J^{ACb} can be further tested. One way would be to run an experiment where participants are asked to watch the kind of clips used above without asking them to type anything explicitly, but instead asking them, after they have watched the clip, to verify one or more fairly specific descriptions of the relevant process-types. The participants responses could either be timed yes/no answers or a ranking of how well the specification fitted howthey thought about the process. The responses could then be tested for correlations across participants. For instance, if participants first were shown the 'see'-clip used above (and thus were likely to report the belief-forming process in terms of 'He saw it'), they could be asked to verify 'He saw it from a distance', 'He saw it behind the basket', 'He saw it in the darkness' etc. The extent to which participants react in similar ways to these queries is an indication of the extent to which they think about the processes in the same way.

J^{ACb} could also be tested by trying to determine whether its presupposition failures are matched by corresponding suspension of judgments of justification. This is trickier to set up though, for it is hard to set up a situation where the participant could go either way

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

concerning how to classify a token. For instance, giving a participant descriptions of a token-process together with the disclaimer that it is uncertain which of these that obtains, mixes up conceptual questions with questions about judgments under uncertainty. Perhaps a better way would be to use developmental studies and see to what extent our capacity for attributing justification and knowledge is tied up with our ability to think about belief-forming processes and our ability to estimate their reliability.

Although all these possible studies are interesting they go beyond the scope of this paper. But before ending the paper I would like to address two possible objections to J^{ACb}, and the interpretation of the data that supports it, in order to preclude some undue criticism.

Objection 1 *(The Existence of the Generality-Problem Controversy):*

> If there is generally convergence in how we type process-tokens in thought, why has it seemed plausible at all that there is a generality problem? The fact that there has been controversy in the literature seems to constitute the premise of a sociological argument against the idea that there is convergence in thought.

The controversy in the literature might be due to a number of factors. These might include, for instance, a tendency in the literature to exaggerate the importance of the idea that a process-token can be logically, or metaphysically typed in a great number of ways, and a tendency to downplay the importance of the fact that it is more natural for us to type a process-token in a certain way. Another reason might be the fact that the demonstration of the generality problem has exclusively been in terms of philosophical counterexamples. These examples are often quite short descriptions of possible situations which are very far from the real-life situations in which we normally attribute justification and knowledge. They leave much to fill in and open up for the possibility of different philosophers filling in in different ways. The situations used in the experiments reported here are, of course, also artificial to some degree, but they are considerably richer when it comes to contextual detail. The distance between philosophical counterexamples and real-life situations might thus have had a non-negligible role in exaggerating the degree of divergence in how we think about process-types in possible situations.

Objection 2 *(Attributor Relativity):*

> J^{AC} makes the process governing a subject's belief depend on the attributor, and too independent of the subject's own thinking. Suppose there is a device in my brain that lets me hear only Sony radios, and I believe p because p is reported on a Sony radio. The canonical description of my cognitive process most people would converge on is probably 'heard on a radio.', and it is likely that this would also be how they think about the process. But surely the process on whose reliability my knowing depends is 'hearing something on a Sony radio' not 'hearing something on a radio'.[30]

This proposed counterexample to K^{AC} (and possibly by extension to J^{AC}) is somewhat misleading. There is no *intuitive* difference in the reliability between hearing something on the radio and hearing something on a Sony radio. And it is this fact that lends credence,

30 I owe this objection to an anonymous reviewer.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

it seems, to the otherwise unsubstantiated claim that an attributor would type the relevant process tokens as hearing something on a radio. However, if there really is no difference in reliability between the two types, the example is not really a counterexample to K^{AC}. K^{AC}, like many forms of reliabilism is immune to 'mistyping' as long as there is no difference in reliability between the alternatives. On the other hand, if there is a difference in reliability between hearing on a radio and hearing on a Sony radio, and the attributor knows this, and she knows about the box, then it seems, quite unlikely that the attributor will type the token-processes in terms of 'heard on a radio'. So, the example is no counterexample in this instance either.

8 CONCLUSION

Empirical evidence has been presented that demonstrates – for a range of normal (reliable and unreliable) belief-forming processes – that people converge on types whose reliability can be used to track justification and knowledge judgments. Several different kinds of reliabilisms that build on these results have been examined, and it has been argued that the reported results are best utilized by J^{ACb}, a form of attributor reliabilism. Even if the specifics of this position are rejected, the very striking correlations reported in Experiment 3 constitute a very strong fundament for reliabilist theories to build upon, and a major empirical confirmation of what reliabilists have been saying for a long time.

ACKNOWLEDGEMENTS

The author would like to thank Kristoffer Ahlstrom-Vij, Staffan Angere, Pernilla Asp, Ruth Chang, Sebastian Enqvist, Alvin Goldman, Simone Jönsson-Asp, Asger Kirkeby-Hinrup, Lisa Miracchi, Erik Olsson, Wlodek Rabinowicz, Stefan Schubert, the participants of the 2012 advanced course on Explicative Epistemology, two anonymous referees and an associate editor for *Episteme* for helpful comments and/or fruitful discussion. The author would also like to thank the kind graduate students who helped out with the piloting of the experiments reported in this paper. This work was made possible by a grant from the Swedish Research Council.

REFERENCES

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

Anglin, J. M. 1977. *Word, Object, and Conceptual Development* New York: W.W. Norton.

Berlin, B., Breedlove, D. E., and Raven, P. H. 1973. "General Principles of Classification and Nomenclature in Folk Biology". *American Anthropologist*, 75(1):214–42.

Brown, R. 1958. "How Shall a Thing be Called?" *Psychological Review*, 65:14–21.

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

Cruse, D. A. 1977. "The Pragmatics of Lexical Specificity". *Journal of Linguistics*, 153–64.

Downing, P. 1980. "Factors Influencing Lexical Choice in Narrative". In W. L. Chafe, ed, *The Pear Stories: Cognitive, Cultural, and Linguistic Aspects of Narrative Production*. Norwood, NJ: Ablex Publishing Corporation.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press

Feldman, R. 1985. "Reliability and Justification". *The Monist*, 68(2):159–75.
Feldman, R. and Conee, E. 2002. "Typing Problems". *Philosophy and Phenomenological Research*, 65(1):98–105.
Goldman, A. I. 1976. "Discrimination and Perceptual Knowledge". *Journal of Philosophy*, 73:771–91.
——. 1979. "What is Justified Belief?" In G. S. Pappas, ed, *Justification and Knowledge,* volume 17 of *Philosophical Studies Series in Philosophy*, 1–23. Dordrecht, Holland: Reidel.
——. 1992. "Epistemic Folkways and Scientific Epistemology". In A. I. Goldman, ed, *Liaisons: Philosophy Meets the Cognitive and Social Sciences*, 155–75. MIT Press/Bradford Books.
Heller, M. 1995. "The Simple Solution to the Problem of Generality". *Noûs*, 29(4):501–15.
John, O. P., Hampson, S. E., and Goldberg, L. R. 1991. "The Basic Level of Personality-trait Hiearchies: Studies of Trait Use and Accessibility in Different Contexts". *Journal of Personality and Social Psychology*, 60:348–61.
Jönsson, M. L. forthcoming. "Linguistic Convergence in Verbs for Belief Forming Processes". *Philosophical Psychology*.
—— (submitted). "The Goodness of Good Categories". Manuscript submitted for publication.
Lewis, D. 1986. "Causal Explanation". In D. Lewis, ed, *Philosophical Papers Volume II*, 214–40. Oxford University Press.
Morris, M. W. and Murphy, G. L. 1990. "Converging Operations on a Basic Level in Event Taxonomies". *Memory and Cognition*, 18(4):407–18.
Murphy, G. L. 2002. *The Big Book of Concepts*. London: MIT Press.
Olsson, E. J. (forthcoming). "A Naturalistic Approach to the Generality Problem". In H. Kornblith and B. McLaughlin, eds, *Goldman and His Critics*. Blackwell.
Pickering, M. and Garrod, S. 2004. "Towards a Mechanistic Psychology of Dialogue". *Behavioural and Brain Sciences*, 27:169–226.
Rosch, E. 1978. "Principles of Categorization". In E. Rosch and B. B. Lloyd, eds, *Cognition and Categorization*. Hillsdale, NJ: Erlbaum.
Rosch, E., Mervis, C. B., Gray, W. D., Johnson, M. D., and Boyes-Braem, P. 1976. "Basic Objects in Natural Categories". *Cognitive Psychology*, 8:382–439.
Rysiew, P. 2011. "Epistemic Contextualism". In E. N. Zalta, ed, *The Stanford Encyclopedia of Philosophy*. Winter 2011 edition.
Tanaka, J. W. and Taylor, M. E. 1991. "Object Categories and Expertise: Is the Basic Level in the Eye of the Beholder". *Cognitive Psychology*, 15:121–49.
Tversky, B. and Hemenway, K. 1983. "Categories of Environmental Scenes". *Cognitive Psychology*, 15:121–49.
Vallacher, R. R. and Wegner, D. M. 2011. "Action Identification Theory". In P. A. M. Van Lange, A. W. Kruglanski, and E. T. Higgins, eds, *Handbook of Theories of Social Psychology*, volume 1, pages 327–49. SAGE Publications Ltd.
Wisniewski, E. J. and Murphy, G. L. 1989. "Superordinate and Basic Category Names in Discourse: A Textual Analysis". *Discourse Processes*, 12:245–61.

MARTIN L. JÖNSSON recently completed a two year Post Doctoral position at the Department of Psychology, City University London, UK, and has now returned to the Department of Philosophy, Lund University, Sweden, where he obtained his Ph.D in 2008. In addition to his work on reliabilist epistemology, he has worked in the philosophy of language on problems for use-based theories of meaning – so far on compositionality, meaning holism and meaning similarity – and in cognitive psychology on issues pertaining to prototype theory, in particular on fallacious reasoning, prototype composition and generic generalizations. His work is published in journals such as *Cognition, Journal of Memory and Language* and *Journal of Philosophical Logic*. Department of Philosophy, Lund University, Kungshuset Lundagård, 22222 Lund, Sweden.

https://doi.org/10.1017/epi.2013.27 Published online by Cambridge University Press