

## Epistemic luck and the generality problem

Kelly Becker

Received: 22 May 2006 / Accepted: 14 May 2007 / Published online: 7 June 2007  
© Springer Science+Business Media B.V. 2007

**Abstract** Epistemic luck has been the focus of much discussion recently. Perhaps the most general knowledge-precluding type is veritic luck, where a belief is true but might easily have been false. Veritic luck has two sources, and so eliminating it requires two distinct conditions for a theory of knowledge. I argue that, when one sets out those conditions properly, a solution to the generality problem for reliabilism emerges.

**Keywords** Reliabilism · Generality problem · Epistemic luck · Veritic luck · Externalism · Theory of knowledge

The various kinds of epistemic luck have been the focus of a lot of attention recently, and for good reason. Any proposed theory of the link between truth and belief constitutive of knowledge must be able to diminish or even eliminate the possibility of merely luckily true beliefs. I will be concerned herein primarily with what has come to be known as veritic epistemic luck, which has two distinct sources. Eliminating veritic luck thus requires two distinct conditions that together provide the requisite belief-truth link constitutive of knowledge: first, a process reliabilist condition construed sufficiently generally to be applicable to more than one belief token, and second, a condition that typically applies to specific belief tokens. It just so happens that this dual condition approach provides resources sufficient to solve the generality problem for reliabilism.

According to reliabilism, knowledge is true belief formed by a reliable process, that is, a process that produces mostly true beliefs, or, better still, one that would produce mostly true beliefs if used many times.<sup>1</sup> The generality problem, about which I shall say more

---

<sup>1</sup> For belief-dependent processes—those at least part of whose inputs are other beliefs—the relevant notion is conditional reliability, where a process is conditionally reliable if it would yield mostly true beliefs when its input beliefs are true. See, for example, Goldman (1979).

below (after motivating my approach to solving it by setting out the conditions necessary to eliminate veritic luck on any adequate theory), is fairly easily grasped. Every belief token is an instance of many process types, and whether a process is deemed reliable depends on how one individuates types. For example, suppose I form a true belief based on a coin toss. One would normally think that this process is unreliable. But one could always cite a more fine-grained individuation of the type to undermine that verdict, say for example forming beliefs about who will win the 2007 NCAA basketball tournament based on flipping this coin on the Monday afternoon before the championship game. If the belief is true, this “process” is correct 100% of the time, hence adequately truth-conducive, hence reliable.<sup>2</sup> So which is the correct individuation of types? Conee and Feldman put the problem succinctly: “Without a specification of the relevant type, process reliabilism is radically incomplete. Only when a bearer of reliability has been identified does the theory have any implications about the [adequacy of the belief-truth link for] beliefs in particular cases” (1998, 3).<sup>3</sup>

## 1 Veritic luck

Since my main concern is veritic luck (and a close cousin of it, on which more below), rather than the very different phenomenon of reflective luck, where the latter presents a problem only for internalist theories of justification, and because my objective is to solve the generality problem for reliabilism, it should be clear that we are working under the general rubric of externalist theories of knowledge.<sup>4</sup> Duncan Pritchard distinguishes the two aforementioned kinds of epistemic luck, where a belief is veritically lucky “if, and only if, the agent’s belief that  $\varphi$  is true in the actual world but false in nearly all near-by

<sup>2</sup> It might *not* be deemed reliable if we were to counterfactualize the reliability condition, as suggested at the beginning of this paragraph. While the process applicable to just this belief is truth-conducive in the actual world, it would presumably not produce mostly true beliefs if applied many times, or throughout close possible worlds, and is thus not reliable. I’ll come back to this suggestion when we return to the generality problem in §II. It’s a useful one, but, I shall argue, implementing it runs the risk of collapsing the distinction between the two anti-luck conditions of our theory, which will leave us with no means to eliminate one of the two kinds of knowledge-precluding luck. I thank an anonymous referee for demanding discussion of the counterfactualization of single case belief-forming mechanisms.

<sup>3</sup> The original has “justification of” where I have placed the bracketed insertion. (See note 4, below, for a partial explanation of the amendment.) So far as I can see, nothing much hangs on this.

<sup>4</sup> One might have noticed that I have just distinguished reflective luck as a problem for internalism about justification from veritic luck as an issue for externalist (and internalist) theories of knowledge. Why the slide from justification to knowledge? The short answer is that, though justification is an important epistemological concept in its own right, once we cast the issue in terms of justification, it threatens to cede too much ground to internalism, where the intuitive concept of justification finds its most natural home. When externalists try to explicate ‘justification’, they often backpedal on their externalist commitments in order to honor inevitably internalist intuitions about the structure of that concept. (Goldman 1979, 1988 are good examples.) I prefer, then, to focus on the belief-truth link constitutive of knowledge, and leave the proper rendering of justification (which, in large part because I am an externalist, I believe is unnecessary for knowledge) to the internalist experts. Some might disagree with all this on the grounds that justification is a general concept, independent of the internalism/externalism debate. But the recent history of epistemological theorizing buttresses my point that, at least very often, justification is thought of as a matter of being in a position to provide reflectively accessible grounds for one’s beliefs. This explains why many believe that one can have justified empirical beliefs even if they are *all* false, for example in the New Evil Demon scenario, where there is *no* belief-truth link, but one has good grounds for one’s beliefs.

possible worlds in which the belief is formed in the same manner as in the actual world” (Pritchard 2003, 111). This is a good starting point, but it needs a small modification. A belief can be veritically lucky when true in the actual world but false in some or many nearby worlds, as opposed to nearly all nearby worlds. Take Zagzebski’s Mary example (Zagzebski 1996, 285). Mary believes that her husband is in the living room based on seeing someone who looks like her husband sitting in her husband’s favorite chair. In fact, her husband’s brother, who looks like her husband, is sitting in the chair, while her husband sits in the corner of the room, unseen by Mary. Her belief is true and formed by a reliable process, but it is veritically lucky. In some nearby worlds, though not almost all of them, Mary’s belief is false—her husband is not in the room but she thinks he is.

In contrast to veritic luck, a belief is reflectively lucky just in case it is true in the actual world but false in nearby worlds *consistent with what one knows only by means that are reflectively accessible*.<sup>5</sup> A stock objection to reliabilism, and to externalism more generally, is that it allows for the compatibility of reflective luck and knowledge. A chicken-sexer, for example, who is somehow reliable in determining the sex of a chick but has no idea how she does it, has reflectively lucky beliefs. That is, given only grounds that are reflectively accessible, her beliefs could easily be false. Cases like this one pull theorists’ intuitions in different directions. Externalists think that the chicken-sexer knows, hence that reflective luck does not preclude knowledge possession, whereas internalists tend to think that the chicken-sexer’s beliefs are unjustified, hence that she does not know, indicating that for them reflective luck does preclude knowledge. A general defense of externalism is beyond the scope of this paper, but since my main aim is to defend reliabilism—generally construed as a brand of externalism—specifically against the generality problem, I will assume an externalist framework herein.

Our first substantive task, then, is to delineate the two sources of veritic luck. A belief can be veritically lucky when formed by a general process that is reliable. The Mary case shows this. Mary’s belief is formed by her truth-conducive vision, but the world “conspires” in a somewhat unusual way to make her belief only luckily true. It’s not that Mary’s general belief-forming process is lucky; rather, it’s reliable. It’s just that her particular belief is luckily true. Let’s call this ‘world luck’, because it’s a peculiar feature of the actual world that makes Mary’s belief veritically lucky. The other strain of veritic luck is instantiated when a true belief is formed from an unreliable process. Suppose an agent believes that *p* is true (where *p* is an ordinary empirical proposition), simply because she wishes that *p*. This is clearly not a reliable (truth-conducive) process, but sometimes beliefs based on wishful thinking are true, hence veritically lucky—in any given case there will be many nearby worlds where the belief is false. Let’s call this ‘process luck’. Later in this section I will describe an important example which shows clearly that the phenomenon of knowledge-precluding process luck can occur even when it does not satisfy the definition of veritic luck, above. (This is the “close cousin” of veritic luck mentioned parenthetically at the very beginning of this section.) It will provide further motivation for my

<sup>5</sup> This formulation of reflective luck is close to Pritchard’s (2006, 17), where he says that reflective luck occurs when an agent has a true belief, but, “given only what the agent is able to know by reflection alone,” it is false in most nearby worlds. But this is ambiguous. It could mean, as I have construed reflective luck, that it is false in nearby worlds consistent with what one knows only by reflectively accessible means, or it could mean that it is false in nearby worlds given only what one knows by reflection. The latter formulation seems too strong, since it would preclude most perceptual beliefs, as they are not based on reflection alone. (It is clear from the context—a discussion of empirical knowledge—that Pritchard does not mean to characterize reflective luck as a problem only, or even primarily, for a priori knowledge.)

dual condition anti-luck epistemology. Whether such examples imply that we should revise the definition of veritic luck or, instead, simply note the fact that there is an additional important strain of process luck for which we need to account in an anti-luck externalist theory is a merely terminological issue. In any event, our theory must avoid the negative implications of process luck.

Since there are two varieties of veritic luck, we'll need two distinct mechanisms for handling them. The obvious fix for process luck is to stipulate that beliefs must be formed from a reliable process to yield knowledge. Wishful thinking doesn't satisfy this condition, but vision might. Of course, vision itself can be unreliable, say in bad lighting or at great distances, and individuating processes in a non-arbitrary way to identify those at issue in determining reliability is the nub of the generality problem. But for now, it suffices to note that vision is generally more reliable than wishful thinking.

What kind of condition do we need to eliminate world luck? Since world luck arises from peculiar features of the world in particular circumstances, we need a belief-truth linking condition that applies to individual instances of belief formation. There are two main candidates in the literature, safety and sensitivity. *Safety* says that a true belief is knowledge only if, in (most or) nearly all of the nearby worlds in which S believes that p, p is true.<sup>6</sup> Mary's belief is unsafe, because there are nearby worlds where she believes that her husband is in the living room, but it is false. These are worlds where her husband's brother is in her husband's chair, but her husband is not in the room. *Sensitivity* says that a true belief is knowledge only if, if p were false, S would not believe that p. Mary's belief is insensitive: If her husband were not in the room, she would believe that he is anyway, mistaking her husband's brother for her husband.

Of the two candidate conditions designed to eliminate world luck, I prefer sensitivity, though it is not very popular any more, primarily because it leads to denying that knowledge is closed under known entailment—denying, that is, the intuitive principle that, *for all agents S and propositions p and q, if S knows that p, and knows that p entails q, then S is in a position to know that q.*<sup>7</sup> I am in the small minority of theorists who believe that the closure principle is false, and so I'm happy to accept sensitivity, but my solution to the generality problem does not require a commitment to sensitivity. Safety will do as well.

When Nozick originally proposed the sensitivity condition, he noted that its application must be relativized to the specific method the agent actually uses to acquire a belief. He illustrated this by his grandmother case (Nozick 1981, 179). Suppose a grandmother believes truly that her grandson is well, based on seeing him. If he were unwell, other family members would tell her that he is fine, and she would believe them. Does this mean she does not know that he is well because, if it were false, she would believe it anyway? It seems not; after all, she *sees* that he is well. Thus a plausible refinement to either sensitivity

<sup>6</sup> Safety is defined and defended by Pritchard (2003, 2005) and Sosa (1999). Pritchard prefers the rendering that reads "nearly all" rather than "most" nearby worlds.

<sup>7</sup> I say here "*is in a position to know that q*" because S may simply fail to draw the relevant inference, hence not know that q. Here's an example of closure failure: S knows that she is sitting at her desk. If it were false, she would not believe it (because she would be standing, lying down, walking, or sitting somewhere else, and thus wouldn't believe she is sitting at her desk). S also knows that, if she is sitting at her desk, then she is not a brain-in-a-vat (BIV). But S does not know that she is not a BIV, because if it were false—if she were a BIV—she would believe she is not a BIV anyway. Nozick (1981) and Dretske (1971) have both defended a version of sensitivity, and both have argued that knowledge is not closed under known entailment.

or safety is to relativize the principles to the actual method the agent uses in forming the specific belief in question. These are the resulting principles:

*Sensitivity:* S knows that p only if, if it were false that p, S would not believe that p, given the method that S actually uses in believing that p.

Comment: Grandmother knows because, if it were false that her grandson is well, she would not believe that he is well, based on her actual method of judging whether he is well by looking at him. (Or: If she were to use her actual method in the closest possible worlds where the grandson is unwell, she would not believe that he is healthy.)

*Safety:* S knows that p only if, in nearly all of the nearby worlds where S believes that p based on the method S actually uses, S's belief is true.

Comment: Grandmother knows because her belief is true in nearly all the nearby worlds where she forms her belief based on her actual method of judging whether her grandson is well by looking at him.

The demand to eliminate world luck, a phenomenon illustrated by the Mary case, requires us to include a modalized tracking principle such as sensitivity or safety in our theory of knowledge. Such principles apply to individual belief tokens and thus utterly lack generality. The grandmother case indicates that any such principle must be relativized to the method of belief formation actually used in a specific instance. The phenomenon of process luck, by contrast, shows that our theory must include a reliable process condition.<sup>8</sup> Eliminating both world luck and process luck requires that we distinguish the *particular methods* by which one forms a belief in a specific instance from the *general processes* from which one forms beliefs. This is the key to solving the generality problem.

Before setting out the generality problem and my proposed solution, I should like to address a thought that may have occurred to the reader. “Why do we need a modalized tracking principle *and* a reliable process condition to eliminate both kinds of veritic luck? Isn't it the case that when a true belief is formed from an unreliable process, it violates the tracking conditions, too? Why not just a tracking theory?” Suppose S forms a lucky true belief that p by wishful thinking. If p were false, S would still believe it. The worry, then, is that process reliability may be redundant—tracking principles can eliminate both world luck and process luck—and so the two-tiered strategy is unmotivated.

Not so fast. Counterexamples can be (and have been) constructed to show that tracking principles alone cannot eliminate knowledge-precluding luck. Suppose Haveit wants me to believe truly that someone in the office owns a Ford.<sup>9</sup> He hypnotizes me and plants the suggestion that Nogot owns a Ford, from which I infer that someone in the office does, when in fact Haveit is the only one in the office who owns a Ford, and if he did not, he would not lead me to believe that anyone in the office does. Thus if it were false that someone in the office owns a Ford, I would not believe that someone does, given my actual specific *method* of believing that someone in the office owns a Ford based on a particular inference subsequent to being hypnotized by Haveit. My belief is therefore sensitive (and safe, see below), but does not constitute knowledge.

<sup>8</sup> Or, if you prefer, a condition that requires true beliefs to be based on general cognitive or intellectual virtues. The early reliabilism of Goldman is sometimes thought of as a precursor to externalist virtue theories. Sosa (1991) and Greco (2000) subsequently shift the focus from reliable processes to stable cognitive or intellectual virtues (or character traits) of the individual. If one were to follow the Sosa/Greco line, one would still face the generality problem in a slightly different guise: How does one non-arbitrarily individuate intellectual virtues or character traits?

<sup>9</sup> This is a modification of an example from Alan Goldman (1987, 184).

How best to individuate methods (rather than processes) is a vexed question whose resolution is beyond the scope of this paper.<sup>10</sup> Following the line of thought initiated by Nozick, Jonathan Vogel provides an exemplar of method individuation which makes clear that the specific content of the belief in question is naturally included in specifying the method of belief formation. Vogel is replying to Goldman's (1983) purported counterexample to Nozick's sensitivity principle. Oscar sees Dach the dachshund and forms the belief *that there is a dog before me*. Dach is standing near a hyena, and Oscar believes falsely that hyenas are dogs, too. Thus in the very close world where Dach is not in sight, Oscar would believe falsely that there is a dog before him, and his belief would violate sensitivity. Vogel's solution is to relativize sensitivity to Oscar's actual method of belief formation: "One way of characterizing Oscar's method would be something like 'If you seem to see a small, squat, elongated animal with short legs and floppy ears, that barks, etc., conclude that there is a dog before you'" (Vogel 1987, 199). In the nearby world where Dach is not in sight, Oscar would not believe that there is a dog before him *by his actual method*, hence his belief does not violate sensitivity.<sup>11</sup> In the Haveit example, my somewhat complicated *process* of forming belief—inference based on hypnotic suggestion—is unreliable. Hypnotic suggestion is a good way of forming false beliefs, for instance that I am the Queen of England, or a dog, or that my feet are on fire.<sup>12</sup> Still, my belief is sensitive, relative to my specific method which, following Vogel's example, would be something like: If you believe that Nogot, who works in your office, owns a Ford, conclude that someone in your office owns a Ford. So not all cases of knowledge-precluding luck can be eliminated by sensitivity (or safety; see below), hence the need for a process reliability condition. And once again, this sort of example compels us to provide criteria for distinguishing specific methods from general processes.

Earlier in this section (see the third paragraph of Section 1), I claimed that there are clear cases of knowledge-precluding epistemic luck that are not cases of veritic luck as we have defined that notion. Haveit may well be one of them (hence my shift from 'veritic' to 'knowledge-precluding' luck when describing the example). In most nearby worlds, my belief regarding whether someone in the office owns a Ford is true—Haveit is hell-bent on making it so. It is this latter fact that makes my belief luckily true. Haveit could cause me to have all kinds of false beliefs, including, as he actually does, the belief that *Nogot* owns a Ford, so there is obviously something unreliable about my belief-forming process (even

<sup>10</sup> Luper-Foy (1984, 29) raises an important issue about method individuation, namely, whether methods should be (1) individuated by reference to the actual causes of belief or (2) given a phenomenalist interpretation. He concludes, tentatively, the latter (at least for perceptual knowledge), which is consistent with Nozick's (1981, 184) claim that methods that are the same "from the inside" are identical. Williamson (2000, §7.4) argues that methods should be individuated externalistically, which (I think) means by reference to external causes. The issues here get complicated pretty fast, but for present purposes it suffices that we individuate methods partly by reference to the *specific content* of the belief in question, which, as I argue in the text, distinguishes methods from processes and thus opens up my strategy for solving the generality problem.

<sup>11</sup> Goldman (1976) also argues for a version of a counterfactual tracking principle for perceptual beliefs. The counterexample in his review article is motivated in part by his view that any such principle should be relativized to perceptual equivalents—a hyena is not a perceptual equivalent of a dachshund. Appeal to methods is a similar (but more general) strategy for avoiding obvious problems with unrelativized tracking principles.

<sup>12</sup> Of course, if Haveit wants to hypnotize me only to instill in me beliefs that he knows are true (or, better still, knows that he knows are true), then this is essentially a non-standard case of reliable testimony. But this is *not* how the case is set up. The whole point is that we are to suspect that Haveit is unreliable. After all, he causes me to believe, in addition to the true proposition that *someone* in the office owns a Ford, the false proposition that *Nogot* owns one.

here, a relatively benign case, it's batting only .500—great for hitting in baseball, not so great for forming beliefs), though the other particular belief issuing from that process—that *someone* in the office owns a Ford—is sensitive. In turn, this shows that externalists must be on guard not only against veritic luck, which sometimes occurs in the form of process luck (other times as world luck), but also against other cases of non-veritic process luck, all the while being permitted, *qua* externalists, to ignore reflective luck.

Now, one is welcome to say that even the Haveit case involves veritic luck, because there may be nearby worlds, depending on how one tells the story, where he does not want me to believe truly that someone in the office owns a Ford. But I recommend that we don't quibble here. By stipulation, in the *closest* possible worlds, Haveit causes me to believe (on false grounds) that someone in the office owns a Ford, and that, in my view anyway, establishes the fact that my belief is *not* veritically lucky—it is true in most of the closest worlds. You may disagree, but most important is that we recognize (1) that process luck and world luck are distinct, (2) that any externalist epistemology must reckon with both kinds of luck, whether all cases of process luck are instances of veritic luck or not, and (3) that modalized tracking principles do not eliminate process luck.

The Haveit case also undermines the safety theorist's claim to eliminate knowledge-precluding luck, because in nearly all the closest worlds where I believe that someone in the office owns a Ford, based on the specific method outlined above, it is true (again because, *ex hypothesi*, Haveit would not lead me to believe someone in the office owns a Ford unless Haveit owned one), and yet the unreliability of hypnosis-based beliefs implies that my belief is only luckily true and surely not an instance of knowledge. Interestingly, there are cases that undermine safety but not sensitivity. Suppose Brown believes that the earth revolves around the sun, based solely on her religious beliefs as a sun-worshipper. In all the nearby worlds where she believes that the earth revolves around the sun, it's true, hence the belief is safe. But basing beliefs on sun-worshipping grounds is unreliable. It leads to many beliefs that are false, for example, that the sun god is angry when the corn fields dry up or that skin cancer is caused by failing to acknowledge the power of the sun god. (By contrast, if it were false that the earth revolves around the sun, Brown would believe it anyway. Her belief is safe but insensitive, which is a nice result for sensitivity theorists.) My paper, "Reliabilism and Safety," gives a more complete defense of the claims that safety does not eliminate knowledge-precluding epistemic luck and that this motivates the inclusion of a process reliability condition.<sup>13</sup>

## 2 The generality problem and its solution

Having motivated the claim that eliminating both kinds of veritic luck, as well as non-veritic, knowledge-precluding process luck, requires a tracking condition *and* a process reliability condition, we can now generate a solution to the generality problem. The

<sup>13</sup> That paper is a response to Pritchard's claim that, once one identifies the right kind of modalized tracking principle, in his view safety, there is no longer a need to supplement the resultant theory with a condition that guarantees reliability: "There is no inherent reason, however, why a safety-based account of knowledge should mention the reliability of processes at all... [T]he appropriate moral to draw... does not seem to be that we need to keep supplementing the reliabilist thesis ad infinitum...but rather that we should simply accept that knowledge is, at root, just true belief that meets the safety principle" (Pritchard 2003, 119). Cases like the sun-worshipper show that eliminating veritic luck, as defined by Pritchard, is not sufficient to preclude all strains of knowledge-precluding epistemic luck. This is why I have taken pains to delineate those strains.

generality problem for reliabilism is that there is no non-circular or non-arbitrary way to individuate processes, and it arises because every token of a belief-forming process belongs to many different types of belief-forming processes.<sup>14</sup> Trying to individuate processes in an intuitively satisfactory and theoretically fruitful way pulls us in opposite directions. Suppose Smith believes truly that there is a bird in the tree, based on her vision in relatively good lighting at a fairly long distance, say 75 m, and the bird happens to be a bright red cardinal. Which general process is at work here? Vision? Vision in good lighting? Vision on Friday afternoons? Vision of brightly colored objects in good lighting at long distance? If we say *vision*, then one might revoke the claim that Smith knows that there is a bird in the tree because vision is not always reliable. It doesn't yield mostly true beliefs about small objects at great distances, or in the dark, or when one is on LSD. So we need to more narrowly individuate the process to get the correct verdict that Smith knows. But cases like this exert pressure to individuate processes so narrowly that they threaten to apply to just one instance of belief formation. For example, in the present case we might say that the process type is vision of brightly colored birds in good lighting from an advantageous perspective at less than 80 m when Smith is neither tired nor intoxicated nor... The more we refine the process, the closer we come to what Richard Feldman has called the single case problem:

If relevant types are characterized very narrowly then the relevant type for some or all process tokens will have only one instance (namely, that token itself). If that token leads to a true belief, then its relevant type is completely reliable, and according to [the proposed individuation criterion], the belief it produces is justified... This is plainly unacceptable, and in the extreme case, where every relevant type has only one instance, [this proposal] has the absurd consequence that all true beliefs are justified and all false beliefs are unjustified. (Feldman 1985, 160–161)

The primary concern here seems to be that, if the process applies to only one instance of belief formation, it is, counterintuitively, either totally reliable or totally unreliable, depending on the truth or falsity of the resulting belief. There is, however, a possible solution to the single case problem, briefly adumbrated above. (See note 2.) We might consider the counterfactual truth-propensity of the process. Let's look at the coin toss example from the beginning of the paper, and let us imagine that the relevant process is forming beliefs about who will win the 2007 NCAA basketball tournament based on flipping this coin on the Monday afternoon before the championship game. If the belief is true, this "process" is correct 100% of the time that it's actually used, which yields the counterintuitive result that the process is reliable. But if one is willing to countenance counterfactual considerations, one could say that the process is not in fact reliable, for had it been used many times, it would not have produced mostly true beliefs; alternatively, it doesn't produce mostly true beliefs throughout close possible worlds. Here the truth-propensity of the process is not determined by its actual success ratio, but by its counterfactual success ratio.

Perhaps the single case problem is not merely the concern that, unless the number of beliefs actually formed by a process is statistically significant, there will be no satisfactory way to determine its reliability just by inspecting the truth ratio of its outputs. As we have just seen, counterfactualizing—considering how the process would perform in close possible worlds—can help with this problem. But the problem, construed this way, arises not only for too narrowly individuated processes; it can arise for any process, no matter how

<sup>14</sup> See, for instance, Feldman (1985), Conee and Feldman (1998), and Feldman and Conee (2002).

widely or narrowly individuated. Suppose Smith just awoke from surgery that repaired a congenital eye defect. The surgery is a success, and for the first time in his life, Smith sees his mother and forms the belief that the person he sees is his mother. Unfortunately, the welling up of emotion sends him into cardiac arrest and he dies. It is not indeterminate whether Smith's vision-based belief-forming process is reliable because there is presumably a fact of the matter about whether it would have produced mostly true beliefs. Since the single case aspect of the generality problem is meant to show that there's something fishy, seemingly ad hoc, about individuating processes so *narrowly* that they appear to be tailored just to fit the case of belief formation at hand, it is doubtful that the nub of the problem concerns how many times a process is *actually* used, since that problem, as we've seen, applies to almost any process, even one as general as vision.

Is the single case problem a problem for any process actually used only once, no matter how individuated, or is it a problem just for those processes so narrowly individuated that they could be used only once? Though the precise nature of the single case problem is not entirely clear, my hunch is that the latter is thought to be the real source of trouble. It is especially difficult to see how a process that is applicable only to one belief token, for example, forming a belief about whether that bird, in this tree, now, is a sparrow, based on seeing it in good lighting at close range, could be *generally* truth-conducive. And yet we've already seen how counterfactualizing can help here, too. There is presumably a fact of the matter about the truth ratio of the outputs of such a process throughout close possible worlds. Perhaps, then, the foe of reliabilism will nonetheless complain that such processes utterly lack generality. If the original inspiration for reliabilism is the thought that there are general cognitive mechanisms that tend to produce true beliefs and that such mechanisms are the objects of a naturalistic inquiry into knowledge, it will be disappointing to find that the mechanisms at issue are different for each belief produced. At this point, the single case problem seems less genuinely problematic than originally advertised. The *principled* objection that it is absurd to think that the reliability of narrow processes which could yield only one belief token is determined by the truth or falsity of that token might be rebutted by going counterfactual. What remains is a concern that such narrowly conceived processes are not really general cognitive mechanisms, and that this seems out of touch with the original motivation for reliabilism.

However, there is a further point to insisting on some level of generality of process individuation, independently of any of the aforementioned concerns. Even if we *could* solve the generality problem by individuating processes in such a manner that they could apply to only one instance of belief formation, we wouldn't want to. We might thereby have a possible solution to the generality problem, but it's not a desirable one because the resulting theory of the belief-truth link constitutive of knowledge would be compatible with cases one would intuitively describe as process luck. Here, at last, we make contact with the overall strategy of this paper. Recall that there are two sources of knowledge-precluding luck: unreliable processes that happen to yield true beliefs in some cases, and quirky features of the actual world that result in an agent's having a reliably formed but still lucky true belief, such as the Mary case. Individuating *processes* so narrowly that they apply to one instance of belief formation is essentially to collapse them into the *methods* discussed in the previous section. Given my method, where methods always attach to belief tokens—they are ways of forming specific beliefs in specific instances—my belief that someone in the office owns a Ford is safe and sensitive, but is nonetheless not a case of knowledge. The clearest diagnosis is that the general belief-forming process is unreliable (that this is a case of process luck), but the diagnosis would not be available to us if processes were as narrowly individuated as methods. That is, if processes were simply

identical to methods, my process *would* be reliable, and no belief-truth linking condition would be violated, even though it is clear both that I don't know that someone in the office owns a Ford, and that my belief is luckily true. Thus to eliminate both kinds of luck, we need to keep our two belief-truth linking conditions distinct—one that applies to individual beliefs and is relativized to method in order to eliminate world luck, and one that applies to processes applicable to more than one belief token. It follows that processes must be individuated somewhat more coarsely than methods.

Seeking to avoid the single case problem pulls us in the other direction, toward individuating processes more generally, but then we confront the no-distinction problem (*ibid.*). The most telling cases occur when an agent forms two (or more) true beliefs from what we would ordinarily deem identical processes, but intuitively one of the beliefs constitutes knowledge and the other does not. Since the reliabilist says that knowledge is true belief formed by a reliable process, and since both true beliefs are formed from the same process, she is committed to saying, contrary to fact, that either both are cases of knowledge, or neither is. Suppose Jones sees a bird in a tree and forms the true beliefs (1) that there is a bird in the tree and (2) that there is a sparrow in the tree.<sup>15</sup> The process at work might be vision within such-and-so distance in such-and-so lighting. Let's stipulate that Jones is very good at detecting the presence of birds, but not as good at determining the type of bird. The problem is that Jones's belief that there is a bird in the tree seems to constitute knowledge,<sup>16</sup> whereas her belief that the bird is a sparrow is veritically lucky. There are lots of nearby worlds where she thinks it's a sparrow, but it's some other kind of small bird. The root of the problem is that, in order to avoid the single case problem (which, I've been arguing, we ought to do whether or not the single case problem is *really* a problem), we seek to individuate processes more generally, and in doing so, cases arise where the processes involved in the production of distinct beliefs are apparently identical, and so identically reliable, but one belief is an instance of knowledge and the other is not.

The reliabilist needs a way to individuate processes that steers a course between the single case problem and the no-distinction problem. My solution, in a nutshell, is to individuate general belief-forming processes very narrowly, but not so narrowly as to land us in the single case "solution", and then let the modalized tracking condition, which is independently motivated by the phenomenon of world luck, eliminate any luck arising from peculiar features of specific cases.

Almost every aspect of my solution to the generality problem has been anticipated by Alvin Goldman. Goldman has argued both for a condition approximating sensitivity (1976) and for process reliabilism.<sup>17</sup> He has also argued for (at least) two distinct means of individuating processes, though so far as I know he has not attempted to combine them. My remaining task is to put all these pieces together.

Goldman (1979, 12) says we should individuate processes content-neutrally, by which he means, I take it, that the specific content of the particular belief in question may not be included in the specification of the process type. For example, if the question is whether Smith knows that *that* (perceived) bird in the tree is a sparrow, we cannot say that the

<sup>15</sup> This is similar to Feldman's (1985) mountain-goat-on-a-distant-hill example.

<sup>16</sup> The generality problem is often construed as a problem for reliabilist theories of *justification*. Once again (see note 4), I aim to be studious in avoiding talk of justification. And as noted before, this is because, for many theorists, the concept of justification is associated with having reflectively accessible grounds for belief, and it's no surprise that externalists will demur from giving an adequate account of that. I prefer, then, to talk generally about the belief-truth link that constitutes knowledge in an externalist framework, and thus to avoid any attempt to elucidate of the concept of justification.

<sup>17</sup> His process reliabilism is found most notably in Goldman (1986), but in many other places as well.

process is *forming belief on the basis of seeing that* perceived bird, now. Adopting content neutrality for process individuation, we escape the single case problem. I'll say more about what I mean exactly by content neutrality presently.

Goldman elsewhere (1986, 50f) claims that the relevant type is the narrowest one that is causally operative in belief production. Conee and Feldman rightly point out that this criterion for process individuation threatens to lead us right back to the single case problem:

Thus, on the reasonable assumption that the content of any normally formed belief is causally determined by its antecedent psychological causes, according to [the narrowest causally operative criterion of process individuation] each relevant type can have only one content for its output belief. This makes trouble in cases in which the proposition believed dictates the truth-ratio of all process types leading only to it. In such cases the reliability of the relevant type is settled by the mere identity of the belief. (Conee and Feldman 1998, 14)

However, in concert with content-neutrality, the single case problem is averted. Together, these criteria entail that the relevant process *type* is the narrowest, *content-neutral* process that is causally operative in belief production.

Content neutrality allows us to avoid individuating processes in an overly particular manner, whereas narrow individuation allows us to circumvent the no-distinction problem. Jones's vision may permit her to be a reliable bird detector, but not a reliable sparrow detector. We can account for this difference by reference to two distinct, content-neutral belief-forming processes: (1) forming beliefs about *birds* on the basis of vision of objects within such-and-so distance in such-and-so lighting, and (2) forming beliefs about *sparrows* on the basis of vision of objects within such-and-so distance in such-and-so lighting. One obvious potential objection is that these processes are not content-neutral. The right thing to say is that, while it is true that we are distinguishing the processes partly by their belief contents, even the sparrow process is *specific* content-neutral. It applies generally to all instances of sparrow belief production, not just to the present case of forming the belief that *that* bird is a sparrow. For such identifying, indexical beliefs, and for beliefs with indexical contents generally, their contents crucially include time and place. Insisting on specific content-neutrality suffices to avoid the single-case problem. Hereafter, then, I will use the expression 'specific content-neutral' when discussing my proposed criterion for process individuation.

Still, one could worry that individuating processes in the narrow manner proposed violates not the letter but the spirit of the single-case problem. Is it really legitimate to distinguish beliefs-about-birds processes from beliefs-about-sparrows processes? Well, first, and as noted, this at least does *not* lead us to the single-case problem *as presented* by Conee and Feldman. On my proposal, the processes are applicable to many instances of belief formation, hence their reliability does not depend essentially on the truth or falsity of a single belief token. Second, such fine distinctions also allow us to correctly differentiate agent-specific reliable processes. Jones is a reliable bird detector, but perhaps not a reliable sparrow detector, whereas Smith, an ornithologist, may also be a reliable sparrow detector in the same conditions. Any sensible epistemology must be able to account for these differences. (I return to this point three paragraphs below.) It is important to note, too, that these fine-grained distinctions are neither *ad hoc* nor arbitrary. I am offering a general recipe for process type individuation that is independent of specific cases.

Let's look briefly at the ultimate test case. Imagine someone, Brown, who reliably forms true beliefs about the presence of one particular bird. Let's say Brown calls the bird 'William'. Does my proposal allow such a unique process? Yes. Brown's process might be

something like this: forming beliefs about the presence of William based on vision in at least dim lighting at close range. This is a general process in the relevant sense, applicable (/repeatable) in many different situations to beliefs with very similar but still distinct contents, given the time and place indexes essential to those contents. Brown's method, on the other hand, is specific content-dependent. If one favors sensitivity as one's world-luck eliminating condition, then Brown's belief constitutes knowledge only if, if it were false that that [bird] is William, Brown would not believe it that it is William, given his method of forming a belief about whether that (there) is (now) William based on Brown's vision in at least dim lighting at close range (and so on).

The present proposal for process type individuation, then, is that the relevant type is the narrowest, specific content-neutral process that is causally operative in belief production. Conee and Feldman present three conditions that any solution to the generality problem must satisfy. The first is that the solution must be principled. It cannot allow for "ad hoc case-by-case selections of types that match our intuitions" about reliability (*ibid.* 4). Content neutrality is specifically designed to rule out ad hoc case-by-case selections, and, relatedly, to avoid the trap of the single case problem.

Second, the solution "must make defensible epistemic classifications... The types identified must have a reliability that is plausibly correlated with the justification of the resulting beliefs" (*ibid.*). Our very narrow but specific content-neutral means of individuation is justified by the fact, briefly noted above, that individuals have very different cognitive abilities, where some, for example, are reliable sparrow detectors and others are not. In order to capture the relevant differences between individual agents, it makes sense to individuate processes very narrowly. Again, it is quite plausible to say that the belief-truth link in Jones's belief that there is a bird in the tree is strong enough to constitute knowledge (hence that the belief is justified, if you insist on talking that way), whereas Jones's belief that there is a sparrow in the tree does not constitute knowledge. Only a narrow process individuation can capture this important difference. Agents other than Jones, however, may be reliable sparrow detectors within such-and-so distance in such-and-so lighting, and for them the belief that there is a sparrow in the tree would constitute knowledge, provided that, in the specific case, world luck is not involved. Notice that narrow process individuation does not eliminate world luck. If there were several fake sparrows in the tree surrounding the real sparrow in whose direction the agent's attention is focused—fakes the agent cannot distinguish from the real one—the agent's true belief would be veritically lucky, an instance of world luck, because neither sensitive nor safe. Sensitivity violation: If it were false that the thing seen is a sparrow, the agent would believe that it is anyway. Safety violation: There are nearby worlds where the agent believes that the thing seen is a sparrow but the belief is false. In this situation, the agent's lack of knowledge is diagnosed as a violation of the modalized tracking condition, which is relativized to the content-specific method of belief formation employed.

Third, "a solution must remain true to the spirit of the reliabilist approach... [It] must somehow implement the basic idea that it is the reliability of a process of belief formation, specified in non-epistemic terms, that settles the epistemic status of the belief" (*ibid.* my emphasis). For example, a version of reliabilism that would violate this requirement might say that a reliable type is one based on adequate evidence, and an unreliable type is one based on inadequate evidence. I submit that the criteria for process individuation presented herein satisfy this condition. Whether a process is reliable, on this account, turns solely on whether the narrowest, causally operative, specific content-neutral type is *truth-conducive*—on whether it produces mostly true beliefs. There is no appeal to justification, adequacy, evidential goodness, or other evaluative terms of epistemic appraisal.

### 3 Concluding remarks

According to the present proposal, then, knowledge is reliably formed true belief that satisfies a modalized tracking principle. The general problem of veritic luck (and its close cousin—process luck that does not involve veritic luck per se) pulls all the strings here. Put differently, the strategy is somewhat top down. We start with the question: What is knowledge? It's true belief formed in a way that involves no epistemic luck. Where does epistemic luck come from? Two sources: unreliable general processes that luckily produce true beliefs and quirky features of the actual world that cause one's reliably formed beliefs to be luckily true. A successful strategy for eliminating luck generates in turn a solution to the generality problem. We must first identify belief-truth linking conditions that eliminate both kinds of luck. One such condition, for example sensitivity, must apply to specific belief tokens to rule out world luck (as in the Mary case), but the condition must also be relativized to the specific method used to be at all plausible (from Nozick's grandmother example). This sort of modal tracking principle was originally designed to apply to specific instances of belief formation. So we let it do its job, and thus no longer have any motivation for individuating processes to fit single cases. But we still need to eliminate process luck. We thus individuate processes very narrowly but specific content-neutrally. The narrowness allows us to account for specific cognitive abilities of agents, perhaps even the ability to recognize a particular bird (or one's own spouse) on various occasions. This solution to the generality problem is *not at all* ad hoc; rather, it is independently motivated by what has been perhaps the most crucial problem for epistemological theories since Gettier first brought it to everyone's attention—epistemic luck.

Finally, it might be complained that the solution to the generality problem presented herein is not really a reliabilist solution at all, insofar as much of the heavy lifting is done by a modalized tracking principle that makes no reference to reliable processes.<sup>18</sup> Even so, the dual condition approach is recognizably reliabilist. What's more, to repeat a main idea from the previous paragraph, process reliabilism *must* be supplemented by some kind of tracking condition (or whatever other condition one might dream up) to eliminate knowledge-precluding world luck, so it's no use complaining that I have so supplemented it.

### References

- Becker, K. (2006). Reliabilism and safety. *Metaphilosophy*, 37(5), 691–704.
- Brown, J. (2004). *Anti-individualism and knowledge*. Cambridge, MA: MIT Press.
- Conee, E., & Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies*, 89, 1–29.
- Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy*, 49(1), 1–22.
- Feldman, R. (1985). Reliability and justification. *The Monist*, 68, 159–174.
- Feldman, R., & Conee, E. (2002). Typing problems. *Philosophy and Phenomenological Research*, 65, 98–105.
- Goldman, Alan (1987). Nozick on knowledge: finding the right connection. In S. Luper-Foy (Ed.), *The possibility of knowledge: Nozick and his critics* (pp. 182–196). Totowa, NJ: Rowman & Littlefield.

<sup>18</sup> An artful dodge here might be to say that, for many, tracking principle-based epistemologies are thought of as versions of reliabilism. “So this is a thoroughgoing reliabilist solution after all.” (See, for example, Goldman (1986, §3.2) and Brown (2004, §4.3), where both distinguish ‘local’ and ‘global’ reliability—the former involving proposition-specific tracking principles and the latter reliable processes. Vogel (2000) refers to Nozickean sensitivity as a kind of reliabilism.) But I'm not going to say that, since it is obvious from the very nature of my project here that I'm working under the assumption that the concept of reliability implies the truth-conduciveness of general belief-forming processes.

- Goldman, Alvin (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73(20), 771–791.
- Goldman, Alvin (1979). What is justified belief? In G. S. Pappas (Ed.), *Justification and knowledge* (pp. 1–23). Dordrecht: D. Reidel.
- Goldman, Alvin (1983). Review of philosophical explanations. *Philosophical Review*, 92(1), 81–88.
- Goldman, Alvin (1986). *Epistemology and cognition*. Cambridge, MA: Harvard UP.
- Goldman, Alvin (1988). Strong and weak justification. *Philosophical Perspectives*, 2, 51–69.
- Greco, J. (2000). *Putting skeptics in their place*. Cambridge: Cambridge University Press.
- Luper-Foy, S. (1984). The epistemic predicament: Knowledge, Nozickian tracking, and scepticism. *Australasian Journal of Philosophy*, 62(1), 26–49.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge, MA: Harvard University Press.
- Pritchard, D. (2003). Virtue epistemology and epistemic luck. *Metaphilosophy*, 34(1/2), 106–130.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford UP.
- Pritchard, D. (2006). Moral and epistemic luck. *Metaphilosophy*, 37(1), 1–25.
- Sosa, E. (1991). *Knowledge in perspective*. Cambridge: Cambridge University Press.
- Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Perspectives*, 13, 141–153.
- Vogel, J. (1987). Tracking, closure, and inductive knowledge. In S. Luper-Foy (Ed.), *The possibility of knowledge: Nozick and his critics* (pp. 197–215). Totowa, NJ: Rowman & Littlefield.
- Vogel, J. (2000). Reliabilism leveled. *The Journal of Philosophy*, 97(11), 602–623.
- Williamson, T. (2000). *Knowledge and its limits*. New York: Oxford University Press.
- Zagzebski, L. (1996). *Virtues of the mind*. Cambridge: Cambridge UP.