RELIABILIST JUSTIFICATION (OR KNOWLEDGE) AS A GOOD TRUTH-RATIO

BY

JONATHAN E. ADLER

Abstract: Fair lotteries offer familiar ways to pose a number of epistemological problems, prominently those of closure and of scepticism. Although these problems apply to many epistemological positions, in this paper I develop a variant of a lottery case to raise a difficulty with the reliabilist's fundamental claim that justification or knowledge is to be analyzed as a high truth-ratio (of the relevant belief-forming processes). In developing the difficulty broader issues are joined including fallibility and the relation of reliability to understanding.

Fair lotteries offer familiar ways to pose a number of epistemological problems, prominently those of closure and of scepticism.[1] Here's one way of posing a sceptical problem: Take your best (purported) example of inductive knowledge. There is some probability, however low, that the evidence is true and the conclusion drawn from it is false. I can specify a fair lottery with enough tickets so that the probability of your losing with a single ticket is greater than the probability that the conclusion of your inductive inference is correct, given the evidence. Yet, you can *know* your conclusion in the inductive case, but not that your ticket will lose in the lottery. How come, given that losing the lottery is more probable?

Answering this question is a problem for a variety of epistemological views, though I use it to raise a difficulty specific to (most versions of) *reliabilism*. However, in developing this difficulty broader issues are joined including fallibilism and the relation of reliability to understanding.

These versions of reliabilism hold that a belief is justified if it arose from a process that tends to produce a preponderance of true rather than

Pacific Philosophical Quarterly 86 (2005) 445–458
© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

false beliefs (a "good truth-ratio"); and it is knowledge, if the resulting belief is undefeated. So, for example, Alston writes,

> The reliability that is in question here [for the identification of justification with reliability or for taking reliability to be an adequate criterion of justification] is the reliability of belief formation and sustenance. To say that a belief was formed in a reliable way is, roughly, to say that it was formed in a way that can be depended upon generally to form true rather than false beliefs. (Alston, 1989, p. 108; later e.g. pp. 232, 244.[2] See also Goldman, 1979, 1986, p. 106, for a rule-based account.)

Reliabilism has been challenged on many counts: for its externalism, for problems in "demon worlds", for the breadth of (possible) worlds covered by "tends" (or "would produce a preponderance . . . if . . ."), for failure to explain the advantage of knowledge over true belief, for permitting maximization of true beliefs to arise accidentally (via cognitive malfunctions), for failure to ground the normativity of epistemic judgments, for the relativity of reliability to attributors, for lacking a principled basis to individuate processes (or to select a privileged reference class), and, perhaps most surprisingly, for "level-confusions". (For sampling of these criticisms, see Plantinga, 1993a, pp. 197–210; Feldman and Conee, 1997; Vogel, 2000. For some defenses see the above references to Goldman, Alston, and also Levin, 1997.) However, as far as I know, the core element of *maximizing true beliefs* has not itself been questioned presumably because justified beliefs purport to satisfy belief's aim of truth non-accidentally, and knowledge does so.

In this paper, I focus on the core element that knowledge, justification, or related epistemic notions are to be analyzed as the reliability of belief-fixation processes and I challenge the explication of reliability as the preponderance or maximization of true beliefs. The initial challenge is posed by the following example, which specifies a lottery-like mechanism directed to the *same* conclusion (belief) as the knowledge-permitting mechanism, unlike the opening example.

A company that manufactures widgets knows that exactly one out of every thousand of their products suffers a singular defect as a by-product of (an ineliminable imperfection in) their excellent – and much better than average – manufacturing system. Whenever there is a defect, it is sufficiently glaring so that customers recognize it and complain. Some managers would like to reduce the percentage of defects. The plan is to introduce a special detector, well designed to read 'OK' just in case the widget is not defective. The detector is to be applied subsequent to the normal manufacturing process to locate any defect before the widgets are shipped to the stores.

Smith and Jones, who both know of the one in a thousand defects, are each given a detector. Batches of widgets are randomly sent to one and

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

later, without either one's knowledge, to the other. (The company will compare their findings afterwards through a check on a proper subset of each of their judgments via a laborious, yet infallible-beyond-a-reasonable-doubt, method.)

As each widget comes to his station, Smith momentarily glances away from the video game he is playing to stamp it 'OK', expressive of his corresponding (degree of) belief, while wholly ignoring his detector. As Smith knows, out of each batch of 1000, he is guaranteed to be correct 999 times by this method (and so better than by use of the detector, as explained next).

Less so Jones. Knowing of her manager's well founded confidence in the detector, Jones applies it to each widget carefully and skillfully, and assigns it 'OK' (mostly) or 'Defect' (rarely) according to her determination. Given the complexity of operating the detector and normal human limits, the probability of an error in any evaluation is .003, though Jones is a first-rate technician. The .003 error rate is a simple average over many shipments of different technicians, applying to both kinds of erroneous judgments (judging an OK widget defective; judging a defective widget OK). The managers know that the detector working properly must give a correct result by the nature of the mechanism, and, in fact, many batches tested with the detector contain no errors at all. Still, enough batches come back with the occasional error to yield the .003 error rate. Attempts to discern the sources of the errors have failed, so it is thought of as random noise resulting from various forms of human error or environmental distortions.

Smith's manager regularly finds him playing video games, and threatens his job. Smith protests that his method yields a truth-ratio that is very high and better (and faster) than that of the esteemed Jones. Smith may add a zinger to his protest: for each widget where his and Jones' judgements differ, he would be willing to *bet* the manager (at even odds) that he is correct. (The laborious process noted above will settle bets.) Still, the manager remains unmoved.

Shorn of the rhetorical device of the manager, I take this example to pose three related challenges for reliabilism.

First, Smith can be more reliable than Jones when he ignores the quality of any widget before him, whereas Jones uses an appropriate and excellent epistemic method of inquiry.[3] (As discussed further below, Smith's method is something like "assigning an object (e.g., a widget) to a category (e.g., 'not defective') on the basis of objects of that kind belonging to the category in an extremely high proportion of cases". So it is lottery-like in assigning at least one member of a set to a category deviant from the rest on the crucial dimension, where the assignment is equally likely to hold of any member of that set. The procedure is otherwise, however, not dependent on a chance mechanism, as with a fair lottery.)

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

Second, Jones can come to *know* (on the basis of her use of the detector) that a widget she tests is OK, but Smith cannot, even though Smith's judgment is more reliable. If knowledge is explicated through subjunctives, then, with a qualification to come, a variant of a standard test, applied to each widget, yields:

(a) Were this widget not OK, Smith would not believe it is.
(b) Were this widget not OK, Jones would not believe it is.[4]

(a) is false, but (b) is true.[5] The falsity of (a) is due to Smith's belief of each widget that it is OK, if he believes this of any. It is not due to a lack of robustness of his method. (Recall Alston's account of reliable processes as ones that "can be depended upon generally to form true rather than false beliefs", which I grant and take to apply to Smith's method.)[6]

Third, Smith can be justified in believing that a widget is non-defective to a certain degree (999/1000), but only Jones can be justified in *all-out* accepting it as true. Put differently, only in Jones' case is there the possibility of *detachment* from a probability assignment to an all-out (unqualified) judgment (of acceptance as true) (Kaplan, 1996, Chs. 3 and 4). Not so for Smith – there is no detachment from the .999 probability that this widget is OK to its all-out (unqualified) assertion.[7,8]

This last challenge responds to a possible reply on behalf of the reliabilist: Smith may not, as implied, (fully or all-out) believe of each widget that it is OK, unlike Jones. (He can, of course, fully believe that the probability that this widget is non-defective is .999.) But this only shifts the locus of the challenge. If Smith will not so believe, it must be because he cannot – his high degree of probability justifies, as he recognizes, no more than a corresponding high degree of partial belief. So the problem now is why Smith cannot fully or all-out believe in cognizance of his method, while Jones can,[9] when Smith is more reliable. It is also an awkward response for the reliabilist to offer, since if a partial belief more accurately reflects the dictates (uncertainties) of his procedure than that is the appropriate – reliable – belief-like attitude.

Like the opening lottery-sceptical problem, the widget-example is set up to pose *comparative* judgments. But, unlike the lottery problem, it is not essentially dependent on comparative judgments, since the lottery-like mechanism is applicable to the same judgment: Smith cannot come to know or to fully believe (accept) that each widget is OK, even if it is.

Still, the comparative problem remains of explaining why Jones is, in central epistemic ways, better off than Smith. The reliabilist seems to require that the explanation be reliability-based and, yet, that it allow for Smith's greater truth-ratio. As the example also indicates, the reliabilist cannot really avoid entering comparative assessments. For it is possible for more than one method to apply to the formation, or ground, of a

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

belief. Each method can yield a preponderance of true beliefs. But they reach conflicting conclusions in the case at hand.[10]

Can the reliabilist bluntly reply that Smith *can* know of each, but one, widget that it is OK? No. Smith's case is designed to be analogous in this respect to a lottery case, and it seems a firm datum that you can not know, in advance, that you will lose a large, fair lottery with one ticket.[11] Corroboration is secured from the data of *assertion*, assuming that knowledge is the proper condition for assertion (Unger, 1975, Ch. VI; DeRose, 1996; Williamson, 2000, Ch. 11). Smith cannot all-out assert of any widget that it is defect-free. His manager would correct him "You mean that you are almost sure that it is defect-free". Jones, however, can so assert.[12] (The point applies even if the reliabilist limits himself to justification. For the relevant notion is that of justification of the proposition believed (sufficient for knowledge). The conditions under which that notion is satisfied are those under which one is entitled to all-out accept the proposition as true (or to assert it without qualification)).

Our reliabilist may persevere, noting that Jones stands a greater chance of error in any case than Smith. What's the difference, the reliabilist continues, between Smith's success in the 999 other cases and Jones' almost certain lesser success? Underlying this proposed reliabilist reply, as in the original lottery-sceptical problem, is a common conflation: between a *fixed probability* (less than one) *of a procedure*, as with the widgets and Smith, and *fallibility*, as with Jones' use of the defective widget detector.[13] The conflation is a kind of levels-confusion: It confuses the (first-order) success of a procedure (or method) for knowing with one's (second-order) confidence in that procedure.

Smith's method guarantees failure, and it is on the basis of that method that each case is judged. However, it is possible for Jones to know of each widget that it is as she judges it, even given the error rate. For it is compatible with a non-extreme probability of one's being mistaken in each of a set of judgments that one is not mistaken in any.[14] Even when she does err, though, knowledge is possible in the remaining cases, radical sceptical responses aside, since the potential errors arise external to her method, as noise or interference.

These differences explain why to introduce fallibility, subsequent to accepting the (justifying) basis for a judgment or belief is just giving with one hand, while taking back with the other. The complete justification or support (for a particular judgment or belief) "screens-off" the impact of one's fallibility, below the appropriate threshold (for complete justification). Above that threshold, the justification simply fails (to be complete).[15]

Early on his recent book, John Hawthorne (2003) assimilates cases of a lawful process with some statistical chance of error (e.g. that Jim, who is

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

young and healthy, will not have a heart attack next week) to lottery cases. Hawthorne backs up the assimilation in a way that raises a problem for the distinction drawn above (or my application of it). Just as we are disinclined to assert all-out "Bill will lose the lottery" (with a single ticket), so too for "Jim will not have a heart attack next week".[16] The worry is that in either case the unwillingness to assert should be taken at face value as indicative of the lack of – impossibility of – knowledge.

This difference, however, is illusory, and it exhibits a give-away/take-back illusion that is especially persuasive with inductive knowledge. We are to assume that there is some stronger claim that you will assert e.g. "Jim will be in Tallahassee next Monday" which you can know. The matter of his not having a heart attack next week is then a (known) entailment of what you do assert, which you have not *specifically* checked. Without engaging the general issue of epistemic closure, I claim that you do already have reasons (evidence) sufficient for the former assertion. Your reluctance is due to the lack of salience of these *background* reasons and the standard *normalcy* assumptions. The lack of salience is partly induced by the assertion, which highlights a familiar (and not way-out) possibility (Jim's having a heart attack) that you have not specifically addressed. However, the suggestion of a counter-possibility is only relevant (as undermining) if it is not already ruled out by your background reasons and normalcy assumptions.

Any inductive inference is susceptible to the same (pragmatic) destabilization by highlighting some of the (future) conditions that must hold for the inference to work (to yield knowledge), and which could not be checked. (In the case at hand, no matter how good your evidence, there is not only the possibility of Jim's having a heart attack, but of his wife taking ill, so he has to cancel his trip, or, of his plane being cancelled, so he does not arrive until Tuesday and so on.) The hesitancy to assert as a datum (to deny knowledge or to assimilate these cases to lottery ones) amounts to giving-away that the issue is not inductive scepticism, and then taking it back. (If Jim arrives in Tallahassee on schedule, along the expected, non-deviant, route, would you think "That was a bit of luck" or "I knew it"? Or, if someone presses you at the original time "Do you know that Jim will not have a heart attack on the journey?", you are hesitant or you outright deny it. Instead, though, try the less misleading "Do you expect conditions to be normal next week in respect of Jim's trip, and if so, will he have a heart attack?". Different response.)

When it is granted that your evidence is sufficient for knowing that Jim will be in Tallahassee on Monday, it is automatically granted – so regularly and so automatically as not to be worth noting – that part of what justifies acceptance on the salient evidence (of, say, Jim's telling you his plans) is the background evidence about Jim's health and other matters related to the real chance of suffering a heart attack and that normal

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

conditions obtain. No separate acts of exclusion of these possibilities are necessary, as is typical of inductive inference.[17] So it is these possibilities which are given-away (as ruled out), when the original claim is conceded as justified. But then these standard concessions are taken-back when they are assimilated to the probability that one will not actually lose the lottery. The assimilation effectively contradicts the concession that conditions are normal. Observe, by contrast, that the tiny chance of winning the lottery is itself a normalcy condition for a fair lottery – the chance that you will win, which is an obstacle to knowing that you will lose, remains. Whatever vagueness and indeterminacy afflicts the notion of normalcy conditions, those conditions are already – independently – assumed in any inductive inference, and that notion draws a line with (pure) lottery cases on the other side. (For a comparison: A deterministic system, with occasional malfunctions, is different in design from an indeterministic system, yielding similar outputs with a very high probability.)

Let me summarize these differences with a simple chart, and if you like, a variant example. Imagine a barrel with a thousand apples. A careful and skilled, but fallible, evaluator examines each, and gets all but two correct as good. Compare her to someone, like Smith, who judges each one good automatically based on the knowledge that 999 are good:

	Fallibility	**High Truth Ratio**
1. Some luck (accident) in being correct from agent's point of view.	Yes	Yes
2. Possibility of getting each correct.	Yes	No
3. Some luck in being correct given the agent's method, including background knowledge and normalcy assumptions.	No	Yes

The illusion is to take the shared affirmative answer for [1], as epistemically crucial (for attribution of knowledge or justification, rather than for betting preferences), and so to warrant the assimilation, whereas it is only [2] and [3] that are crucial.

At this juncture I am prepared to introduce a procedure that constitutes an alternative to those of both Smith and Jones. Brown engages in an examination like Jones, but then factors in the 1/1000 rate of defects as a prior probability, via a version of Bayes' Theorem. Let H = This widget is OK, E = The device reads 'OK'.

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

$$\begin{aligned}\Pr(H/E) &= \text{pr}(E/H)\text{pr}(H)/[\text{pr}(E/H)\text{pr}(H) + \text{pr}(E/{-}H)\text{pr}({-}H)] \\ &= (.997)(.999)/[(.997)(.999) + (.003)(.001)] \\ &= .996003/.996006 \\ &= .999997\end{aligned}$$

So Brown's judgments are more probably correct than that of either Smith or Jones. Does Brown provide the reliabilist with a reply to the original problem, since Brown's judgment is more reliable than Smith's and, in part, just because he, like Jones, does examine each widget?

I think not. Earlier, I cast doubt on treating Jones' error rate as a likelihood ratio, as I cast doubt that Smith's .001 probability of misjudging is a kind of fallibility. The crucial question for Brown arises when there is a discrepancy between his individual judgment (like Jones) and the high probability that the widget is OK (i.e. when he judges that a widget is defective): Does he – is he entitled to – understand his examination as adequate to all-out accept (detach or assert) that the widget is defective or, rather, as retaining some probability, however small, of error? If the former, then he is in Jones' position; if the latter, then Smith's, even though, in either case, he has greater assurance. Consequently, the reliabilist cannot appeal to Brown's method as a reply, since ultimately Brown's position just parallels that of either Smith or Jones.[18]

Still, the reliabilist may offer a different way (than the appeal to Brown's procedure) to protest that the account of Smith's method is too simple. A consequence of Smith's mechanical reliance on his knowledge of the 1 in a 1000 defects is that he has a *poorer understanding* than Jones of the quality of the widgets that he certifies as OK, since Jones studies each widget and judges in accord with a proper application of an appropriate, though fallible, method. The reliabilist may now press this observation by the question: Why cannot this difference in understanding itself be incorporated into the description of the process Smith uses, as contrasted to Jones? Alternatively, even if the process remains as described, does not the poorer understanding that Smith suffers entail a worse overall truth-ratio than Jones?

The former suggestion lacks any intuitive grounding. The relation between understanding and judgments of truth seems to be neither lawlike nor robust (in an environment of greater carelessness, Smith's method, even with its diminished understanding, would be epistemically preferable). The proposal also seems in conflict with the reliabilist's core externalism, since to further Smith's understanding requires his grasp of, and inferences from, the success (and failures) of his procedures and methods.

Adding to the description of the process that Smith uses something to the effect that it reflects his understanding of the quality of the widgets does not correspond to any normal or natural process of belief-formation.

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

The *type of process* that Smith uses is, as mentioned above, something like "assigning an object (e.g., a widget) to a category (e.g., 'not defective') on the basis of objects of that kind belonging to the category in an extremely high proportion of cases". This is how we would describe the process, and absent the comparisons I draw (with Jones) this is exactly how we would explain why Smith is so successful (reliable). It captures the full resources of Smith's method without any irrelevancies (e.g., "and entering the assignment in English on an iMac"). Further, since it is avowedly Smith's method, we can refer to his judgment on the matter. Smith's judgment is that his method is as above specified – it guarantees 999 out of each 1000 are correct, and that is better (i.e. more reliable than) any other available method like Jones'.[19]

The question of what type of process to assign to a particular instance introduces the vexed "Generality Problem". But, as has recently been argued, this problem should be viewed as roughly the old problem of fixing the correct *reference class* (Beebe, 2004).[20] Again simplifying, the right reference class type to assign a specific (token) process is "homogenous". It has no partition, whereby membership in this narrower class is (statistically) relevant to membership in the type.[21] Smith's method as we specified it above meets this condition. There is no such partitioning – no further relevant way to break up the process so as to "screen-off " Smith's method, rendering it (statistically) irrelevant.

The alternative suggestion is that the difference in understanding between Smith and Jones is to be cashed out as a difference in their respective truth-ratios. Presumably, the result would be favorable to Jones, and so it would undermine our use of the example to criticize reliabilism. As stated above, the description of the process remains the same. However, the claim now is that the truth-ratio should be determined not only directly by those beliefs generated by the process, and whose content is that a certain widget is OK [defective]. It should extend, as well, to the processes' *indirect* consequences for belief: the intellectual and inferential import for the truth-ratio of one's corpus of beliefs due to the use of, and learning from, this process.[22]

But this result seems not to fit the spirit of reliabilism either, since we are extending the consequences to evaluate the relevant processes beyond that of its standard domain of functioning. The proposal is also fraught with contingency and indeterminacy, even if it is granted that poorer understanding is correlated with a worse truth-ratio. Whether the poorer understanding yields a worse truth-ratio depends upon the happenstance of the frequency and representativeness of the samples to which it is applied, as well as whether the agent tries to compensate for his poorer understanding by practice and study. Further, we think a poor or mediocre understanding is intrinsically (epistemically) bad, not merely consequentially so.[23]

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

Even if the expected differences in understanding cannot aid the reliabilist, these differences do highlight an open path to handle the central problems posed here. As this open path (of according epistemic priority to understanding) might be developed, it is different from the main approaches, since seeking to dislodge knowledge either as a candidate for analysis or as of central epistemological concern.[24] But, of course, less radical deviations remain viable. Broadly, the problems posed reflect a loss in the shift from causal (or causal-explanatory) theories to reliability theories: the loss of the demand for an appropriate connection between the causal type of one's believing and that which renders true the proposition believed. To be specific: The criticisms do not (prima facie) apply to versions of reliabilism that require a probability of 1 for knowledge (or justification or acceptance) or to forms of reliabilism or externalism or subjunctive analyses that do not explicate their key notion via good (or threshold) truth-ratios (e.g. Armstrong, 1971, Ch. 12).[25] Internalist views that demand conclusive reasons are not threatened: Smith can not know because that a widget is a member of a set with a very high probability of belonging to a category is not a conclusive reason to assign it to that category.[26,27]

Department of Philosophy
Brooklyn College and the Graduate School, CUNY

NOTES

[1] Hawthorne (2003) shows that lottery cases can be used to discuss a wide range of epistemological problems.

[2] Alston favors (1989, p. 244) a reliable (indication-of-a) ground, rather than a reliable process, view. The difference does not bear on the challenges posed below.

[3] As already noted, Smith's method is better for some epistemic purposes, especially betting behavior as a reflection of degrees of belief, and it is better on a reliabilist construal of both Jamesian epistemic goals: seeking truth and avoiding error. See Riggs, 2003.

[4] The condition derives from Nozick, 1981, Ch. 3 section I (and the differences apply to his extension of the analysis to methods). See also Dretske, 1970.

[5] In the nearest worlds where the widget is defective, Jones so detects it, even though of each widget there is a small chance that her judgment is in error.

[6] If it is held that it remains somewhat accidental that Smith is right, when he is, this just restates a problem I am pressing for reliabilism: Why should this make a difference if the truth-ratio remains high? Moreover, how can this result be avoided if a probability less than 1 of a correct belief is compatible with reliabilist justification or knowledge?

[7] This difference allows for a simple response to the opening lottery-sceptical problem: the lottery does not permit, as normal inductive inferences do, all-out acceptance of the conclusion as true. The impossibility of detachment in the lottery case is one way of resolving the "lottery paradox": Of each ticket, a holder of it can not all-out accept that he will lose (prior to selection of a winner). For a recent discussion see Nelkin, 2000.

[8] This example differs from familiar cases of reliable belief formation via an unreasonable to believe process (e.g. Bonjour, 1985, p. 38). Smith's basis is not unreasonable. A consequence

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

is that this example does not suffer the flaw that the familiar cases do: if the belief-formation is via a first-personally unreasonable to believe process, like clairvoyance, it is not credible that the agent actually does all-out believe its output.

[9] If knowledge and fallibility are compatible, Jones can first-personally recognize her (limited) fallibility, psychological barriers aside.

[10] In his 1999, Goldman understands "veritistic epistemology" as "concerned with the production of knowledge, where knowledge is here understood in the 'weak' sense of *true belief*". (5) Later, he writes:

> Under veritism we are asked to select the social practices that would best advance the cause of knowledge (p. 79).

He then offers an extended example (p. 81) of weather forecasting, which explicitly involves comparative claims.

[11] Cohen (1998) denies this datum. He holds that we misjudge otherwise because the mention of the defect rate makes it *salient* that this is a *chance* matter, generating a context in which the attribution of knowledge fails. But in the original context (prior to the mention), where standards are appropriately low, Smith would know (if his belief is not defeated). I think Cohen rightly perceives that the contextualist must bite this bullet. Still, it is an intuitive and theoretically defensible (e.g. in application to the lottery paradox (note 7 above)) datum that in no context, however low the standards, can one know that one will lose. Also, and in addition to the objections in the text, first, this reply is an attempt to explain-away our ordinary judgments, whereas contextualists stress the support that they receive from our ordinary judgments and second, the maneuver does not work for Smith. The judgement that Smith cannot know does not reside in a chance process, like a lottery.

[12] DeRose (1996) presents a problematic case similar to one posed by Harman, 1986, p. 71. In the *newspaper-lottery* case there is about a one in a million chance of getting a defective newspaper which has transposed the scores of a basketball game (in which the Bulls won). With Alice learning the scores from the newspaper, DeRose holds that we judge:

a. Alice can know [assert] that the Bulls won.

But not:

b. Alice can know [assert] that her newspaper is not the bad one.

That we so judge is explained by the SCA (subjunctive conditional analysis), since it is true that:

a.′ Were it not true that the Bulls won, Alice would not believe it.

But not

b.′ Were it not true that Alice's newspaper is not the bad one, Alice would not believe it.

Does the SCA have the implications DeRose assumes? Among the nearest or most similar worlds in which the antecedent of a′ holds are not just worlds in which the Bulls lost and her error-free newspaper so reports (so Alice would not believe that they won). There will also be worlds in which she does have the bad newspaper. In the latter case, it does still report that the Bulls won, and she does believe it. So (a′) is false as well. The realization of the very improbable event, though it leads to a great divergence in the future, does not

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

itself constitute a difference in similarity of worlds by a standard set of criteria. (For those criteria see Lewis, 1986.)

[13] Compare to e.g. Cohen, 1988, who identifies fallibility with compatibility of knowledge and the non-elimination of a contrary to what is known. See Adler, 2002, §10.3.

[14] Compare to Dretske's (1981, Ch. 5) distinction between the information a source carries and the channel on which the delivery of this information depends. (See also his 1970).

[15] As suggested above, a betting test to measure strength of belief obscures this distinction. So far as wagering goes, all that matters is the probability of error, regardless of the source.

[16] A Gricean difference applies, though I won't rest with it. To assert that "Jim will not have a heart attack next week" is informative only if his not having the heart attack next week is special or particularly germane to the context at hand. But we are to imagine it is neither – it is a claim that follows merely from other assertions about Jim e.g. "Jim will be in Tallahassee next week". Thus, it is likely to be taken to implicate that the speaker has some special knowledge about Jim's health, and if not, he should not assert it.

[17] In his discussion, Hawthorne (2003) rightly credits Vogel (1990) with noticing the wide reach of statistical and lottery cases. But Hawthorne does not take sufficient notice of Vogel's correlative claim (in defending epistemic closure) of the pivotal role of background knowledge or evidence.

[18] For more realistic Smith-Jones-Brown comparisons: Assume that you are refereeing a paper for a journal that has a 98% rejection rate. For your epistemic, rather than strictly professional, duty, should you follow Smith or Jones or Brown? Or, think of airplane security guards adopting methods that are the analogues of those of Smith, Jones and Brown, respectively. In these comparisons, it is assumed that there are no indirect costs e.g., that the one using a method akin to Smith's is not recognized by others as so judging, since that would introduce a spreading corrupting factor.

[19] What if Smith understands his method via the broader reference class of "assigning an object to a category on the basis of objects of that kind belonging to the category in a majority or higher proportion of cases"? On this version, his method would not have a better overall truth-ratio than Jones, even though it it is more successful in this application.

[20] Actually, Beebe's solution is two-fold with the reference class determination only one part. The other part of his solution does not, however, bear on the issues here.

[21] Beebe rightly takes the reliabilist to require that the homogeneity of the reference class be *objective*, rather than epistemic or subjective (p. 190). But *if* an objective determination can not make use of background information on what properties or predicates are appropriate for partitioning, how will it avoid gerrymandered or "grue"-like predicates from being forever available? See here Salmon, 1989, pp. 68–83, whom Beebe explicitly follows. A related matter: Beebe eliminates clearly irrelevant properties, e.g. the irrelevance to the reliability of a perceptual process that it takes place on Wednesdays, as not statistically relevant. But more challenging cases are those in which a certain property is statistically relevant, but irrelevant to the lawfulness of the process. We can easily conceive that males and females, a property which Beebe treats like Wednesdays, do have statistically relevant differences in perceptual reliability as a consequence of differences in (social) roles, and so the objects that they attend to, and the conditions under which they so attend. The deeper irrelevance is brought out by finding some further property e.g. a "common cause", to affect another partition under which these (accidentally) statistically relevant properties get screened-off. But are the screening-off properties fixed objectively or, instead, do they require reference to our background knowledge, as well as to lawfulness, which is also, arguably, not purely objective?

[22] Alternatively, the suggestion can be expressed as a variant of the former one which seeks to re-describe the process in a more exclusionary and complex way as something like '. . . when an alternative process is available that allows for judging each case on its merits'. Some of

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

my critical comments in the text below apply, with suitable modifications, to this proposal. But its main defect is that if, as intended, this alternatively described process yields that Smith's method is not superior to Jones', it does so only by conceding that "judging each case on its merits" is epistemically advantageous to a method that yields a higher truth-ratio.

[23] Although the criticism applies to the use of the process as a method to judge, rather than restricting it to an analysis of knowledge or justification, the problem raised is realistic for both. Psychological studies show that our intuitive physics is (largely) Aristotelian rather than Newtonian; and (in some ways) Newtonian rather than Einsteinian (e.g., as to whether time moves uniformly with external conditions (of motion) irrelevant.) Within our environment of (comparatively) slow speeds and weak gravitational forces, conditions are not generated that would falsify our intuitive physics. So our by-and-large high truth-ratio judgments are correlative to a poorer understanding (of the underlying physics). Or, consider the disfavored role of *base rates* conversationally. Speakers generally offer routine descriptions of individuals in a way that is stereotypical for a certain category, so that their assertions, or its presuppositions, will be by-and-large correct (the person so described is actually a member of that category.) So within a normal conversational context – assuming that we are predominantly cooperative speakers (truthful, relevant, informative, brief), who intend to successfully communicate, representativeness is a truth-maximizing way to judge. However, to judge probability or frequency on that basis expresses, and, presumably, encourages, a poor understanding (in the experimental applications, specifically, it overrides regression and the law of large numbers) (Kahneman and Tversky, 1982). The problem for the reliabilist is that he must then take beliefs formed via the representativeness heuristic as justified (or, when true, knowledge), and this will not be due to envisaging the process as either fragile or to operate in way-out, though 'nice', demon world. (See Goldman 1986: 107 and 1988). Goldman actually does allow that use of the representativeness heuristic could be admissible in a right J-rule system: "Perhaps the R-routine [representativeness-routine] is *generally* quite reliable, though it breeds errors in this subclass of cases" (Goldman, 1986, p. 321). But his full discussion in sections 14.2 and 14.4 should be consulted as they contain qualifications, especially to handle multiple systems. See also the discussion of wishful thinking or Bonjour examples in Goldman, 1979 and Levin, 1997. It is doubtful whether there is any nearby possible world, where we have a robust conversational practice and speakers do not provide mostly true (and truth implicating) assertions.

[24] For the former, without sympathy for any replacement thesis, see Williamson, 2000, Ch. 1; for the latter, and with sympathy, see Elgin, 1996, Ch. 4.

[25] Plantinga's "proper function account" seems susceptible to the problems raised here, since the examples require no defect in faculties or in environments or in the relevant processes operating in accord with their design plan. His fourth – avowedly reliabilist – condition is that "there is a high statistical or objective probability that a belief produced . . . is true" (1993b, p. 194; also: Ch. 1, section III).

[26] However, internalist or partially internalist views that explicate justification as a good truth-ratio or unlikelihood of error are threatened e.g. Alston, 1989, p. 84; Haack, 1993, p. 203.

[27] My thanks to Peter Graham and Michael Levin for suggestions and discussion. I am extremely grateful to a referee for *Pacific Philosophical Quarterly* for extensive and valuable comments.

REFERENCES

Adler, J. E. (2002). *Belief's Own Ethics*. Cambridge, MA: The MIT Press.
Alston, W. P. (1989). *Epistemic Justification*. Ithaca, NY: Cornell University Press.

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.

Armstrong, D. M. (1973). *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.

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

Bonjour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.

Cohen, S. (1988). "How to be a Fallibilist," *Philosophical Perspectives* 2, pp. 581–605.

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

DeRose, K. (1996). "Knowledge, Assertion and Lotteries," *Australasian Journal of Philosophy* 74, pp. 568–579.

Dretske, F. (1970). "Epistemic Operators," *The Journal of Philosophy* 69, pp. 1015–1016.

Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: The MIT Press.

Elgin, C. Z. (1996). *Considered Judgment*. Princeton, NJ: Princeton University Press.

Feldman, R. and Conee, E. (1997). "The Generality Problem for Reliabilism," *Philosophical Studies* XCIII, pp. 1–29.

Goldman, A. I. (1979). "What is Justified Belief?," in G. Pappas (ed.), *Justification and Knowledge*. Dordrecht: Reidel, pp. 1–24.

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

Goldman, A. I. (1988). "Strong and Weak Justification," *Philosophical Perspectives, 2: Epistemology*, J. Tomberlin (ed.). Atascadero, CA: Ridgeview, pp. 51–69.

Goldman, A. I. (1999). *Knowledge in a Social World*. Oxford: Oxford University Press.

Haack, S. (1993). *Evidence and Inquiry*. Oxford, Blackwell.

Harman, G. (1986). *Change in View*. Cambridge, MA: The MIT Press.

Hawthorne, J. (2003). *Knowledge and Lotteries*. New York, Oxford University Press.

Kahneman, D. and Tversky, A. (1982). "Subjective Probability: A Judgement of Representativeness" in D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press, pp. 32–47.

Kaplan, M. (1996). *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.

Levin, M. (1997). "You Can Always Count on Reliabilism," *Philosophy and Phenomenological Research* LVII, pp. 607–613.

Lewis, D. (1986). "Counterfactual Dependence and Time's Arrow," in his *Philosophical Papers Vol.II*, Oxford: Oxford University Press, pp. 32–66.

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

Nelkin, D. K. (2000). "The Lottery Paradox, Knowledge, and Rationality," *The Philosophical Review* 109, pp. 373–409.

Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

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

Plantinga, A. (1993b). *Warrant and Proper Function*. Oxford: Oxford University Press.

Riggs, W. D. (2003). "Balancing our Epistemic Goals," *Noûs* 37, pp. 342–352.

Unger, P. (1975). *Ignorance*. Oxford: Oxford University Press.

Vogel, J. (1990). "Are There Counterexamples to the Closure Principle?," in M. D. Roth and G. Ross (eds.), *Doubting*. Dordrecht: Kluwer, pp. 13–27.

Vogel, J. (2000). "Reliabilism Leveled," *The Journal of Philosophy* XCVII, pp. 602–623.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

© 2005 The Author
Journal compilation © 2005 University of Southern California and Blackwell Publishing Ltd.