PHILOSOPHICAL STUDIES SERIES
IN PHILOSOPHY

VOLUME 17

JUSTIFICATION AND KNOWLEDGE

New Studies in Epistemology

Edited by

GEORGE S. PAPPAS
The Ohio State University

D. REIDEL PUBLISHING COMPANY
DORDRECHT : HOLLAND / BOSTON : U.S.A.
LONDON : ENGLAND

ALVIN I. GOLDMAN

WHAT IS JUSTIFIED BELIEF?

The aim of this paper is to sketch a theory of justified belief. What I have in mind is an explanatory theory, one that explains in a general way why certain beliefs are counted as justified and others as unjustified. Unlike some traditional approaches, I do not try to prescribe standards for justification that differ from, or improve upon, our ordinary standards. I merely try to explicate the ordinary standards, which are, I believe, quite different from those of many classical, e.g., 'Cartesian', accounts.

Many epistemologists have been interested in justification because of its presumed close relationship to knowledge. This relationship is intended to be preserved in the conception of justified belief presented here. In previous papers on knowledge,[1] I have denied that justification is necessary for knowing, but there I had in mind 'Cartesian' accounts of justification. On the account of justified belief suggested here, it *is* necessary for knowing, and closely related to it.

The term 'justified', I presume, is an evaluative term, a term of appraisal. Any correct definition or synonym of it would also feature evaluative terms. I assume that such definitions or synonyms might be given, but I am not interested in them. I want a set of ***substantive*** conditions that specify when a belief is justified. Compare the moral term 'right'. This might be defined in other ethical terms or phrases, a task appropriate to meta-ethics. The task of normative ethics, by contrast, is to state substantive conditions for the rightness of actions. Normative ethics tries to specify non-ethical conditions that determine when an action is right. A familiar example is act-utilitarianism, which says an action is right if and only if it produces, or would produce, at least as much net happiness as any alternative open to the agent. These necessary and sufficient conditions clearly involve no ethical notions. Analogously, I want a theory of justified belief to specify in non-epistemic terms when a belief is justified. This is not the only kind of theory of justifiedness one might seek, but it is one important kind of theory and the kind sought here.

In order to avoid epistemic terms in our theory, we must know which terms are epistemic. Obviously, an exhaustive list cannot be given, but here are some examples: 'justified', 'warranted', 'has (good) grounds', 'has reason

George S. Pappas (ed.), Justification and Knowledge, 1–23.

(to believe)', 'knows that', 'sees that', 'apprehends that', 'is probable' (in an epistemic or inductive sense), 'shows that', 'establishes that', and 'ascertains that'. By contrast, here are some sample non-epistemic expressions: 'believes that', 'is true', 'causes', 'it is necessary that', 'implies', 'is deducible from', and 'is probable' (either in the frequency sense or the propensity sense). In general, (purely) doxastic, metaphysical, modal, semantic, or syntactic expressions are not epistemic.

There is another constraint I wish to place on a theory of justified belief, in addition to the constraint that it be couched in non-epistemic language. Since I seek an explanatory theory, i.e., one that clarifies the underlying source of justificational status, it is not enough for a theory to state 'correct' necessary and sufficient conditions. Its conditions must also be appropriately deep or revelatory. Suppose, for example, that the following sufficient condition of justified belief is offered: 'If S senses redly at t and S believes at t that he is sensing redly, then S's belief at t that he is sensing redly is justified.' This is not the kind of principle I seek; for, even if it is correct, it leaves unexplained *why* a person who senses redly and believes that he does, believes this justifiably. Not every state is such that if one is in it and believes one is in it, this belief is justified. What is distinctive about the state of sensing redly, or 'phenomenal' states in general? A theory of justified belief of the kind I seek must answer this question, and hence it must be couched at a suitably deep, general, or abstract level.

A few introductory words about my *explicandum* are appropriate at this juncture. It is often assumed that whenever a person has a justified belief, he knows that it is justified and knows what the justification is. It is further assumed that the person can state or explain what his justification is. On this view, a justification is an argument, defense, or set of reasons that can be given in support of a belief. Thus, one studies the nature of justified belief by considering what a person might *say* if asked to defend, or justify, his belief. I make none of these sorts of assumptions here. I leave it an open question whether, when a belief *is* justified, the believer *knows* it is justified. I also leave it an open question whether, when a belief is justified, the believer can *state* or *give* a justification for it. I do not even assume that when a belief is justified there is something 'possessed' by the believer which can be called a 'justification'. I do assume that a justified belief gets its status of being justified from some processes or properties that make it justified. In short, there must be some justification-conferring processes or properties. But this does not imply that there must be an argument, or reason, or anything else, 'possessed' at the time of belief by the believer.

I

A theory of justified belief will be a set of principles that specify truth-conditions for the schema ⌜S's belief in p at time t is justified⌝, i.e., conditions for the satisfaction of this schema in all possible cases. It will be convenient to formulate candidate theories in a recursive or inductive format, which would include (A) one or more base clauses, (B) a set of recursive clauses (possibly null), and (C) a closure clause. In such a format, it is permissible for the predicate 'is a justified belief' to appear in recursive clauses. But neither this predicate, nor any other epistemic predicate, may appear in (the antecedent of) any base clause.[2]

Before turning to my own theory, I want to survey some other possible approaches to justified belief. Identification of problems associated with other attempts will provide some motivation for the theory I shall offer. Obviously, I cannot examine all, or even very many, alternative attempts. But a few sample attempts will be instructive.

Let us concentrate on the attempt to formulate one or more adequate base-clause principles.[3] Here is a classical candidate:

(1) If S believes p at t, and p is indubitable for S (at t), then S's belief in p at t is justified.

To evaluate this principle, we need to know what 'indubitable' means. It can be understood in at least two ways. First, 'p is indubitable for S' might mean: 'S has no *grounds* for doubting p'. Since 'ground' is an epistemic term, however, principle (1) would be inadmissible on this reading, for epistemic terms may not legitimately appear in the antecedent of a base-clause. A second interpretation would avoid this difficulty. One might interpret 'p is indubitable for S' psychologically, i.e., as meaning 'S is psychologically incapable of doubting p'. This would make principle (1) admissible, but would it be correct? Surely not. A religious fanatic may be psychologically incapable of doubting the tenets of his faith, but that doesn't make his belief in them justified. Similarly, during the Watergate affair, someone may have been so blinded by the aura of the Presidency that even after the most damaging evidence against Nixon had emerged he was still incapable of doubting Nixon's veracity. It doesn't follow that his belief in Nixon's veracity was justified.

A second candidate base-clause principle is this:

(2) If S believes p at t, and p is self-evident, then S's belief in p at t is justified.

To evaluate this principle, we again need an interpretation of its crucial term, in this case 'self-evident'. On one standard reading, 'evident' is a synonym for

'justified'. '*Self*-evident' would therefore mean something like 'directly justified', 'intuitively justified', or 'non-derivatively justified'. On this reading 'self-evident' is an epistemic phrase, and principle (2) would be disqualified as a base-clause principle.

However, there are other possible readings of '*p* is self-evident' on which it isn't an epistemic phrase. One such reading is: 'It is impossible to understand *p* without believing it'.[4] According to this interpretation, trivial analytic and logical truths might turn out to be self-evident. Hence, any belief in such a truth would be a justified belief, according to (2).

What does 'it is *impossible* to understand *p* without believing it' mean? Does it mean '*humanly* impossible'? That reading would probably make (2) an unacceptable principle. There may well be propositions which humans have an innate and irrepressible disposition to believe, e.g., 'Some events have causes'. But it seems unlikely that people's inability to refrain from believing such a proposition makes every belief in it justified.

Should we then understand 'impossible' to mean 'impossible in principle', or 'logically impossible'? If that is the reading given, I suspect that (2) is a vacuous principle. I doubt that even trivial logical or analytic truths will satisfy this definition of 'self-evident'. Any proposition, we may assume, has two or more components that are somehow organized or juxtaposed. To understand the proposition one must 'grasp' the components and their juxtaposition. Now in the case of *complex* logical truths, there are (human) psychological operations that suffice to grasp the components and their juxtaposition but do not suffice to produce a belief that the proposition is true. But can't we at least *conceive* of an analogous set of psychological operations even for simple logical truths, operations which perhaps are not in the repertoire of human cognizers but which might be in the repertoire of some conceivable beings? That is, can't we conceive of psychological operations that would suffice to grasp the components and componential-juxtaposition of these simple propositions but do not suffice to produce *belief* in the propositions? I think we can conceive of such operations. Hence, for any proposition you choose, it will be possible for it to be understood without being believed.

Finally, even if we set these two objections aside, we must note that self-evidence can at best confer justificational status on relatively few beliefs, and the only plausible group are beliefs in necessary truths. Thus, other base-clause principles will be needed to explain the justificational status of beliefs in contingent propositions.

The notion of a base-clause principle is naturally associated with the idea

of 'direct' justifiedness, and in the realm of contingent propositions first-person-current-mental-state propositions have often been assigned this role. In Chisholm's terminology, this conception is expressed by the notion of a '*self-presenting*' state or proposition. The sentence 'I am thinking', for example, expresses a self-presenting proposition. (At least I shall *call* this sort of content a 'proposition', though it only has a truth value given some assignment of a subject who utters or entertains the content and a time of entertaining.) When such a proposition is true for person *S* at time *t*, *S* is justified in believing it at *t*: in Chisholm's terminology, the proposition is 'evident' for *S* at *t*. This suggests the following base-clause principle.

(3) If *p* is a self-presenting proposition, and *p* is true for *S* at *t*, and *S* believes *p* at *t*, then *S*'s belief in *p* at *t* is justified.

What, exactly, does 'self-presenting' mean? In the second edition of *Theory of Knowledge*, Chisholm offers this definition: "*h* is self-presenting for *S* at *t* =df. *h* is true at *t*; and necessarily, if *h* is true at *t*, then *h* is evident for *S* at *t*."[5] Unfortunately, since 'evident' is an epistemic term, 'self-presenting' also becomes an epistemic term on this definition, thereby disqualifying (3) as a legitimate base-clause. Some other definition of self-presentingness must be offered if (3) is to be a suitable base-clause principle.

Another definition of self-presentation readily comes to mind. 'Self-presentation' is an approximate synonym of 'self-intimation', and a proposition may be said to be self-intimating if and only if whenever it is true of a person that person believes it. More precisely, we may give the following definition.

(SP) Proposition *p* is self-presenting if and only if: necessarily, for any *S* and any *t*, if *p* is true for *S* at *t*, then *S* believes *p* at *t*.

On this definition, 'self-presenting' is clearly not an epistemic predicate, so (3) would be an admissible principle. Moreover, there is initial plausibility in the suggestion that it is *this* feature of first-person-current-mental-state propositions – viz., their truth guarantees their being believed – that makes beliefs in them justified.

Employing this definition of self-presentation, is principle (3) correct? This cannot be decided until we define self-presentation more precisely. Since the operator 'necessarily' can be read in different ways, there are different forms of self-presentation and correspondingly different versions of principle (3). Let us focus on two of these readings: a '*nomological*' reading and a '*logical*' reading. Consider first the nomological reading. On this definition a proposition is self-presenting just in case it is nomologically necessary that if *p* is true for *S* at *t*, then *S* believes *p* at *t*.[6]

Is the nomological version of principle (3) – call it '(3_N)' – correct? Not at

all. We can imagine cases in which the antecedent of (3_N) is satisfied but we would not say that the belief is justified. Suppose, for example, that p is the proposition expressed by the sentence 'I am in brain-state B', where 'B' is shorthand for a certain highly specific neural state description. Further suppose it is a nomological truth that anyone in brain-state B will ipso facto *believe* he is in brain-state B. In other words, imagine that an occurrent belief with the content 'I am in brain-state B' is realized whenever one is in brain-state B.[7] According to (3_N), any such belief is justified. But that is clearly false. We can readily imagine circumstances in which a person goes into brain-state B and therefore has the belief in question, though this belief is by no means justified. For example, we can imagine that a brain-surgeon operating on S artificially induces brain-state B. This results, phenomenologically, in S's suddenly believing – out of the blue – that he is in brain-state B, without any relevant antecedent beliefs. We would hardly say, in such a case, that S's belief that he is in brain-state B is justified.

Let us turn next to the logical version of (3) – call it '(3_L)' – in which a proposition is defined as self-presenting just in case it is logically necessary that if p is true for S at t, then S believes p at t. This stronger version of principle (3) might seem more promising. In fact, however, it is no more successful than (3_N). Let p be the proposition 'I am awake' and assume that it is logically necessary that if this proposition is true for some person S and time t, then S believes p at t. This assumption is consistent with the further assumption that S frequently believes p when it is false, e.g., when he is dreaming. Under these circumstances, we would hardly accept the contention that S's belief in this proposition is always justified. But nor should we accept the contention that the belief is justified when it is *true*. The truth of the proposition logically guarantees that the belief is *held*, but why should it guarantee that the belief is *justified*?

The foregoing criticism suggests that we have things backwards. The idea of self-presentation is that truth guarantees belief. This fails to confer justification because it is compatible with there being belief without truth. So what seems necessary – or at least sufficient – for justification is that belief should guarantee truth. Such a notion has usually gone under the label of '*infallibility*', or '*incorrigibility*'. It may be defined as follows.

(INC) Proposition p is incorrigible if and only if: necessarily, for any S and any t, if S believes p at t, then p is true for S at t.

Using the notion of incorrigibility, we may propose principle (4).

(4) If p is an incorrigible proposition, and S believes p at t, then S's belief in p at t is justified.

As was true of self-presentation, there are different varieties of incorrigibility, corresponding to different interpretations of 'necessarily'. Accordingly, we have different versions of principle (4). Once again, let us concentrate on a nomological and a logical version, (4_N) and (4_L) respectively.

We can easily construct a counterexample to (4_N) along the lines of the belief-state/brain-state counterexample that refuted (3_N). Suppose it is nomologically necessary that if anyone believes he is in brain-state B then it is true that he is in brain-state B, for the only way this belief-state is realized is through brain-state B itself. It follows that ' I am in brain-state B' is a nomologically incorrigible proposition. Therefore, according to (4_N), whenever anyone believes this proposition at any time, that belief is justified. But we may again construct a brain-surgeon example in which someone comes to have such a belief but the belief isn't justified.

Apart from this counterexample, the general point is this. Why should the fact that S's believing p guarantees the truth of p imply that S's belief is justified? The nature of the guarantee might be wholly fortuitous, as the belief-state/brain-state example is intended to illustrate. To appreciate the point, consider the following related possibility. A person's mental structure might be such that whenever he believes that p will be true (of him) a split second later, then p is true (of him) a split second later. This is because, we may suppose, his believing it brings it about. But surely we would not be compelled in such a circumstance to say that a belief of this sort is justified. So why should the fact that S's believing p guarantees the truth of *p precisely at the time of belief* imply that the belief is justified? There is no intuitive plausibility in this supposition.

The notion of *logical* incorrigibility has a more honored place in the history of conceptions of justification. But even principle (4_L), I believe, suffers from defects similar to those of (4_N). The mere fact that belief in p logically guarantees its truth does not confer justificational status on such a belief.

The first difficulty with (4_L) arises from logical or mathematical truths. Any true proposition of logic or mathematics is logically necessary. Hence, any such proposition p is logically incorrigible, since it is logically necessary that, for any S and any t, if S believes p at t then p is true (for S at t). Now assume that Nelson believes a certain very complex mathematical truth at time t. Since such a proposition is logically incorrigible, (4_L) implies that Nelson's belief in this truth at t is justified. But we may easily suppose that this belief of Nelson is not at all the result of proper mathematical reasoning, or even the result of appeal to trustworthy authority. Perhaps Nelson believes

this complex truth because of utterly confused reasoning, or because of hasty and ill-founded conjecture. Then his belief is not justified, contrary to what (4_L) implies.

The case of logical or mathematical truths is admittedly peculiar, since the truth of these propositions is assured independently of any beliefs. It might seem, therefore, that we can better capture the idea of 'belief logically guaranteeing truth' in cases where the propositions in question are *contingent*. With this in mind, we might restrict (4_L) to *contingent* incorrigible propositions. Even this amendment cannot save (4_L), however, since there are counterexamples to it involving purely contingent propositions.

Suppose that Humperdink has been studying logic – or, rather, pseudo-logic – from Elmer Fraud, whom Humperdink has no reason to trust as a logician. Fraud has enunciated the principle that any disjunctive proposition consisting of at least 40 distinct disjuncts is very probably true. Humperdink now encounters the proposition p, a contingent proposition with 40 disjuncts, the 7th disjunct being 'I exist'. Although Humperdink grasps the proposition fully, he doesn't notice that it is entailed by 'I exist'. Rather, he is struck by the fact that it falls under the disjunction rule Fraud has enunciated (a rule I assume Humperdink is not *justified* in believing). Bearing this rule in mind, Humperdink forms a belief in p. Now notice that p is logically incorrigible. It is logically necessary that if anyone believes p, then p is true (of him at that time). This simply follows from the fact that, first, a person's believing anything entails that he exists, and second, 'I exist' entails p. Since p is logically incorrigible, principle (4_L) implies that Humperdink's belief in p is justified. But surely, given our example, that conclusion is false. Humperdink's belief in p is not at all justified.

One thing that goes wrong in this example is that while Humperdink's belief in p logically implies its truth, Humperdink doesn't *recognize* that his believing it implies its truth. This might move a theorist to revise (4_L) by adding the requirement that S 'recognize' that p is logically incorrigible. But this, of course, won't do. The term 'recognize' is obviously an epistemic term, so the suggested revision of (4_L) would result in an inadmissible base-clause.

II

Let us try to diagnose what has gone wrong with these attempts to produce an acceptable base-clause principle. Notice that each of the foregoing attempts confers the status of 'justified' on a belief without restriction on *why* the belief is held, i.e., on what *causally initiates* the belief or *causally sustains* it.

The logical versions of principles (3) and (4), for example, clearly place no restriction on causes of belief. The same is true of the nomological versions of (3) and (4), since nomological requirements can be satisfied by simultaneity or cross-sectional laws, as illustrated by our brain-state/belief-state examples. I suggest that the absence of causal requirements accounts for the failure of the foregoing principles. Many of our counterexamples are ones in which the belief is caused in some strange or unacceptable way, e.g., by the accidental movement of a brain-surgeon's hand, by reliance on an illicit, pseudo-logical principle, or by the blinding aura of the Presidency. In general, a strategy for defeating a noncausal principle of justifiedness is to find a case in which the principle's antecedent is satisfied but the belief is caused by some faulty belief-forming process. The faultiness of the belief-forming process will incline us, intuitively, to regard the belief as unjustified. Thus, correct principles of justified belief must be principles that make causal requirements, where 'cause' is construed broadly to include sustainers as well as initiators of belief (i.e., processes that determine, or help to overdetermine, a belief's continuing to be held.)[8]

The need for causal requirements is not restricted to base-clause principles. Recursive principles will also need a causal component. One might initially suppose that the following is a good recursive principle: 'If S justifiably believes q at t, and q entails p, and S believes p at t, then S's belief in p at t is justified'. But this principle is unacceptable. S's belief in p doesn't receive justificational status simply from the fact that p is entailed by q and S justifiably believes q. If what causes S to believe p at t is entirely different, S's belief in p may well not be justified. Nor can the situation be remedied by adding to the antecedent the condition that S justifiably believes that q entails p. Even if he believes this, and believes q as well, he might not put these beliefs together. He might believe p as a result of some other wholly extraneous, considerations. So once again, conditions that fail to require appropriate causes of a belief don't guarantee justifiedness.

Granted that principles of justified belief must make reference to causes of belief, what kinds of causes confer justifiedness? We can gain insight into this problem by reviewing some faulty processes of belief-formation, i.e., processes whose belief-outputs would be classed as unjustified. Here are some examples: confused reasoning, wishful thinking, reliance on emotional attachment, mere hunch or guesswork, and hasty generalization. What do these faulty processes have in common? They share the feature of *unreliability*: they tend to produce *error* a large proportion of the time. By contrast, which species of belief-forming (or belief-sustaining) processes are intuitively

justification-conferring? They include standard perceptual processes, remembering, good reasoning, and introspection. What these processes seem to have in common is *reliability*: the beliefs they produce are generally true. My positive proposal, then, is this. The justificational status of a belief is a function of the reliability of the process or processes that cause it, where (as a first approximation) reliability consists in the tendency of a process to produce beliefs that are true rather than false.

To test this thesis further, notice that justifiedness is not a purely categorical concept, although I treat it here as categorical in the interest of simplicity. We can and do regard certain beliefs as more justified than others. Furthermore, our intuitions of comparative justifiedness go along with our beliefs about the comparative reliability of the belief-causing processes.

Consider perceptual beliefs. Suppose Jones believes he has just seen a mountain-goat. Our assessment of the belief's justifiedness is determined by whether he caught a brief glimpse of the creature at a great distance, or whether he had a good look at the thing only 30 yards away. His belief in the latter sort of case is (*ceteris paribus*) more justified than in the former sort of case. And, if his belief is true, we are more prepared to say he *knows* in the latter case than in the former. The difference between the two cases seems to be this. Visual beliefs formed from brief and hasty scanning, or where the perceptual object is a long distance off, tend to be wrong more often than visual beliefs formed from detailed and leisurely scanning, or where the object is in reasonable proximity. In short, the visual processes in the former category are less reliable than those in the latter category. A similar point holds for memory beliefs. A belief that results from a hazy and indistinct memory impression is counted as less justified than a belief that arises from a distinct memory impression, and our inclination to classify those beliefs as '*knowledge*' varies in the same way. Again, the reason is associated with the comparative reliability of the processes. Hazy and indistinct memory impressions are generally less reliable indicators of what actually happened; so beliefs formed from such impressions are less likely to be true than beliefs formed from distinct impressions. Further, consider beliefs based on inference from observed samples. A belief about a population that is based on random sampling, or on instances that exhibit great variety, is intuitively more justified than a belief based on biased sampling, or on instances from a narrow sector of the population. Again, the degree of justifiedness seems to be a function of reliability. Inferences based on random or varied samples will tend to produce less error or inaccuracy than inferences based on non-random or non-varied samples.

Returning to a categorical concept of justifiedness, we might ask just *how* reliable a belief-forming process must be in order that its resultant beliefs be justified. A precise answer to this question should not be expected. Our conception of justification is *vague* in this respect. It does seem clear, however, that *perfect* reliability isn't required. Belief-forming processes that *sometimes* produce error still confer justification. It follows that there can be justified beliefs that are false.

I have characterized justification-conferring processes as ones that have a 'tendency' to produce beliefs that are true rather than false. The term 'tendency' could refer either to *actual* long-run frequency, or to a 'propensity', i.e., outcomes that would occur in merely *possible* realizations of the process. Which of these is intended? Unfortunately, I think our ordinary conception of justifiedness is vague on this dimension too. For the most part, we simply assume that the 'observed' frequency of truth versus error would be approximately replicated in the actual long-run, and also in relevant counterfactual situations, i.e., ones that are highly 'realistic', or conform closely to the circumstances of the actual world. Since we ordinarily assume these frequencies to be roughly the same, we make no concerted effort to distinguish them. Since the purpose of my present theorizing is to capture our ordinary conception of justifiedness, and since our ordinary conception is vague on this matter, it is appropriate to leave the theory vague in the same respect.

We need to say more about the notion of a belief-forming '*process*'. Let us mean by a 'process' a *functional operation* or procedure, i.e., something that generates a *mapping* from certain states – 'inputs' – into other states – 'outputs'. The outputs in the present case are states of believing this or that proposition at a given moment. On this interpretation, a process is a *type* as opposed to a *token*. This is fully appropriate, since it is only types that have statistical properties such as producing truth 80% of the time; and it is precisely such statistical properties that determine the reliability of a process. Of course, we also want to speak of a process as *causing* a belief, and it looks as if types are incapable of being causes. But when we say that a belief is caused by a given process, understood as a functional procedure, we may interpret this to mean that it is caused by the particular *inputs* to the process (and by the intervening events 'through which' the functional procedure carries the inputs into the output) on the occasion in question.

What are some examples of belief-forming 'processes' construed as functional operations? One example is reasoning processes, where the inputs include antecedent beliefs and entertained hypotheses. Another example is functional procedures whose inputs include desires, hopes, or emotional states

of various sorts (together with antecedent beliefs). A third example is a memory process, which takes as input beliefs or experiences at an earlier time and generates as output beliefs at a later time. For example, a memory process might take as input a belief *at* t_1 that Lincoln was born in 1809 and generate as output a belief *at* t_n that Lincoln was born in 1809. A fourth example is perceptual processes. Here it isn't clear whether inputs should include states of the environment, such as the distance of the stimulus from the cognizer, or only events within or on the surface of the organism, e.g., receptor stimulations. I shall return to this point in a moment.

A critical problem concerning our analysis is the degree of generality of the process-types in question. Input–output relations can be specified very broadly or very narrowly, and the degree of generality will partly determine the degree of reliability. A process-type might be selected so narrowly that only one instance of it ever occurs, and hence the type is either completely reliable or completely unreliable. (This assumes that reliability is a function of *actual* frequency only.) If such narrow process-types were selected, beliefs that are intuitively unjustified might be said to result from perfectly reliable processes; and beliefs that are intuitively justified might be said result from perfectly unreliable processes.

It is clear that our ordinary thought about process-types slices them broadly, but I cannot at present give a precise explication of our intuitive principles. One plausible suggestion, though, is that the relevant processes are *content-neutral*. It might be argued, for example, that the process of *inferring p whenever the Pope asserts p* could pose problems for our theory. If the Pope is infallible, this process will be perfectly reliable; yet we would not regard the belief-outputs of this process as justified. The content-neutral restriction would avert this difficulty. If relevant processes are required to admit as input beliefs (or other states) with *any* content, the aforementioned process will not count, for its input beliefs have a restricted propositional content, viz., '*the Pope* asserts *p*'.

In addition to the problem of 'generality' or 'abstractness' there is the previously mentioned problem of the '*extent*' of belief-forming processes. Clearly, the causal ancestry of beliefs often includes events outside the organism. Are such events to be included among the 'inputs' of belief-forming processes? Or should we restrict the extent of belief-forming processes to '*cognitive*' events, i.e., events within the organism's nervous system? I shall choose the latter course, though with some hesitation. My general grounds for this decision are roughly as follows. Justifiedness seems to be a function of how a cognizer deals with his environmental input, i.e., with the goodness or badness of the operations that register and transform the stimulation that reaches him. ('Deal with', of course, does not mean *purposeful* action; nor is it restricted to *conscious* activity.) A justified belief is, roughly speaking, one that results from cognitive operations that are, generally speaking, good or successful. But '*cognitive*' operations are most plausibly construed as operations of the cognitive faculties, i.e., 'information-processing' equipment *internal* to the organism.

With these points in mind, we may now advance the following base-clause principle for justified belief.

(5) If *S*'s believing *p* at *t* results from a reliable cognitive belief-forming process (or set of processes), then *S*'s belief in *p* at *t* is justified.

Since 'reliable belief-forming process' has been defined in terms of such notions as belief, truth, statistical frequency, and the like, it is not an epistemic term. Hence, (5) is an admissible base-clause.

It might seem as if (5) promises to be not only a successful base clause, but the only principle needed whatever, apart from a closure clause. In other words, it might seem as if it is a necessary as well as a sufficient condition of justifiedness that a belief be produced by reliable cognitive belief-forming processes. But this is not quite correct, give our provisional definition of 'reliability'.

Our provisional definition implies that a reasoning process is reliable only if it generally produces beliefs that are true, and similarly, that a memory process is reliable only if it generally yields beliefs that are true. But these requirements are too strong. A reasoning procedure cannot be expected to produce true belief if it is is applied to false premisses. And memory cannot be expected to yield a true belief if the original belief it attempts to retain is false. What we need for reasoning and memory, then, is a notion of '*conditional reliability*'. A process is conditionally reliable when a sufficient proportion of its output-beliefs are true *given that its input-beliefs are true*.

With this point in mind, let us distinguish *belief-dependent* and *belief-independent* cognitive processes. The former are processes *some* of whose inputs are belief-states.[9] The latter are processes *none* of whose inputs are belief-states. We may then replace principle (5) with the following two principles, the first a base-clause principle and the second a recursive-clause principle.

(6_A) If *S*'s belief in *p* at *t* results ('immediately') from a belief-independent process that is (unconditionally) reliable, then *S*'s belief in *p at t* is justified.

(6B) If S's belief in p at t results ("immediately") from a belief-dependent process that is (at least) conditionally reliable, and if the beliefs (if any) on which this process operates in producing S's belief in p at t are themselves justified, then S's belief in p at t is justified.[10]

If we add to (6A) and (6B) the standard closure clause, we have a complete theory of justified belief. The theory says, in effect, that a belief is justified if and only it is '*well-formed*', i.e., it has an ancestry of reliable and/or conditionally reliable cognitive operations. (Since a dated belief may be over-determined, it may have a number of distinct ancestral trees. These need not all be full of reliable or conditionally reliable processes. But at least one ancestral tree must have reliable or conditionally reliable processes throughout.)

The theory of justified belief proposed here, then, is an *Historical* or *Genetic* theory. It contrasts with the dominant approach to justified belief, an approach that generates what we may call (borrowing a phrase from Robert Nozick) '*Current Time-Slice*' theories. A Current Time-Slice theory makes the justificational status of a belief wholly a function of what is true of the cognizer *at the time* of belief. An Historical theory makes the justificational status of a belief depend on its prior history. Since my Historical theory emphasizes the reliability of the belief-generating processes, it may be called '*Historical Reliabilism*'.

The most obvious examples of Current Time-Slice theories are 'Cartesian' Foundationalist theories, which trace all justificational status (at least of contingent propositions) to current mental states. The usual varieties of Coherence theories, however, are equally Current Time-Slice views, since they too make the justificational status of a belief wholly a function of *current* states of affairs. For Coherence theories, however, these current states include all other beliefs of the cognizer, which would not be considered relevant by Cartesian Foundationalism. Have there been other Historical theories of justified belief? Among contemporary writers, Quine and Popper have Historical epistemologies, though the notion of 'justification' is not their avowed *explicandum*. Among historical writers, it might seem that Locke and Hume had Genetic theories of sorts. But I think that their Genetic theories were only theories of ideas, not of knowledge or justification. Plato's theory of recollection, however, is a good example of a Genetic theory of knowing.[11] And it might be argued that Hegel and Dewey had Genetic epistemologies (if Hegel can be said to have had a clear epistemology at all).

The theory articulated by (6A) and (6B) might be viewed as a kind of 'Foundationalism,' because of its recursive structure. I have no objection

to this label, as long as one keeps in mind how different this 'diachronic' form of Foundationalism is from Cartesian, or other 'synchronic' varieties of, Foundationalism.

Current Time-Slice theories characteristically assume that the justificational status of a belief is something which the cognizer is able to know or determine at the time of belief. This is made explicit, for example, by Chisholm.[12] The Historical theory I endorse makes no such assumption. There are many facts about a cognizer to which he lacks 'privileged access', and I regard the justificational status of his beliefs as one of those things. This is not to say that a cognizer is necessarily ignorant, at any given moment, of the justificational status of his current beliefs. It is only to deny that he necessarily has, or can get, knowledge or true belief about this status. Just as a person can know without knowing that he knows, so he can have justified belief without knowing that it is justified (or believing justifiably that it is justified.)

A characteristic case in which a belief is justified though the cognizer doesn't know that it's justified is where the original evidence for the belief has long since been forgotten. If the original evidence was compelling, the cognizer's original belief may have been justified; and this justificational status may have been preserved through memory. But since the cognizer no longer remembers how or why he came to believe, he may not know that the belief is justified. If asked now to justify his belief, he may be at a loss. Still, the belief *is* justified, though the cognizer can't demonstrate or establish this.

The Historical theory of justified belief I advocate is connected in spirit with the causal theory of knowing I have presented elsewhere.[13] I had this in mind when I remarked near the outset of the paper that my theory of justified belief makes justifiedness come out closely related to knowledge. Justified beliefs, like pieces of knowledge, have appropriate histories; but they may fail to be knowledge either because they are false or because they founder on some other requirement for knowing of the kind discussed in the post-Gettier knowledge-trade.

There is a variant of the Historical conception of justified belief that is worth mentioning in this context. It may be introduced as follows. Suppose S has a set B of beliefs at time t_0, and some of these beliefs are *un*justified. Between t_0 and t_1 he reasons from the entire set B to the conclusion p, which he then accepts at t_1. The reasoning procedure he uses is a very sound one, i.e., one that is conditionally reliable. There is a sense or respect in which we are tempted to say that S's belief in p at t_1 is 'justified'. At any rate, it is tempting to say that the *person* is justified in believing p at t. Relative to his antecedent cognitive state, he did as well as could be expected: the *transition*

from his cognitive state at t_0 to his cognitive state at t_1 was entirely sound. Although we may acknowledge this brand of justifiedness – it might be called '*Terminal-Phase Reliabilism*' – it is not a kind of justifiedness so closely related to knowing. For a person to know proposition p, it is not enough that the *final phase* of the process that leads to his belief in p be sound. It is also necessary that some entire history of the process be sound (i.e., reliable or conditionally reliable).

Let us return now to the Historical theory. In the next section of the paper, I shall adduce reasons for strengthening it a bit. Before looking at these reasons, however, I wish to review two quite different objections to the theory.

First, a critic might argue that *some* justified beliefs do not derive their justificational status from their causal ancestry. In particular, it might be argued that beliefs about one's current phenomenal states and intuitive beliefs about elementary logical or conceptual relationships do not derive their justificational status in this way. I am not persuaded by either of these examples. Introspection, I believe, should be regarded as a form of retrospection. Thus, a justified belief that I am 'now' in pain gets its justificational status from a relevant, though brief, causal history.[14] The apprehension of logical or conceptual relationships is also a cognitive process that occupies time. The psychological process of 'seeing' or 'intuiting' a simple logical truth is very fast, and we cannot introspectively dissect it into constituent parts. Nonetheless, there are mental operations going on, just as there are mental operations that occur in *idiots savants*, who are unable to report the computational processes they in fact employ.

A second objection to Historical Reliabilism focuses on the reliability element rather than the causal or historical element. Since the theory is intended to cover all possible cases, it seems to imply that for any cognitive process C, if C is reliable in possible world W, then any belief in W that results from C is justified. But doesn't this permit easy counterexamples? Surely we can imagine a possible world in which wishful thinking is reliable. We can imagine a possible world where a benevolent demon so arranges things that beliefs formed by wishful thinking usually come true. This would make wishful thinking a reliable process in that possible world, but surely we don't want to regard beliefs that result from wishful thinking as justified.

There are several possible ways to respond to this case and I am unsure which response is best, partly because my own intuitions (and those of other people I have consulted) are not entirely clear. One possibility is to say that in the possible world imagined, beliefs that result from wishful thinking *are*

justified. In other words we reject the claim that wishful thinking could never, intuitively, confer justifiedness.[15]

However, for those who feel that wishful thinking couldn't confer justifiedness, even in the world imagined, there are two ways out. First, it may be suggested that the proper criterion of justifiedness is the propensity of a process to generate beliefs that are true *in a non-manipulated environment*, i.e., an environment in which there is no purposeful arrangement of the world either to accord or conflict with the beliefs that are formed. In other words, the suitability of a belief-forming process is only a function of its success in '*natural*' situations, not situations of the sort involving benevolent or malevolent demons, or any other such manipulative creatures. If we reformulate the theory to include this qualification, the counterexample in question will be averted.

Alternatively, we may reformulate our theory, or reinterpret it, as follows. Instead of construing the theory as saying that a belief in possible world W is justified if and only if it results from a cognitive process that is reliable in W, we may construe it as saying that a belief in possible world W is justified if and only if it results from a cognitive process that is reliable *in our world*. In short, our conception of justifiedness is derived as follows. We note certain cognitive processes in the actual world, and form beliefs about which of these are reliable. The ones we believe to be reliable are then regarded as justification-conferring processes. In reflecting on hypothetical beliefs, we deem them justified if and only if they result from processes already picked out as justification-conferring, or processes very similar to those. Since wishful thinking is not among these processes, a belief formed in a possible world W by wishful thinking would not be deemed justified, even if wishful thinking is reliable *in* W. I am not sure that this is a correct reconstruction of our intuitive conceptual scheme, but it would accommodate the benevolent demon case, at least if the proper thing to say in that case is that the wishful-thinking-caused beliefs are unjustified

Even if we adopt this strategy, however, a problem still remains. Suppose that wishful thinking turns out to be reliable *in the actual world*![16] This might be because, unbeknownst to us at present, there is a benevolent demon who, lazy until now, will shortly start arranging things so that our wishes come true. The long-run performance of wishful thinking will be very good, and hence even the new construal of the theory will imply that beliefs resulting from wishful thinking (in *our* world) are justified. Yet this surely contravenes our intuitive judgment on the matter.

Perhaps the moral of the case is that the standard format of a 'conceptual

analysis' has its shortcomings. Let me depart from that format and try to give a better rendering of our aim and the theory that tries to achieve that aim. What we really want is an *explanation* of why we count, or would count, certain beliefs as justified and others as unjustified. Such an explanation must refer to our *beliefs* about reliability, not to the actual *facts*. The reason we *count* beliefs as justified is that they are formed by what we *believe* to be reliable belief-forming processes. Our beliefs about which belief-forming processes are reliable may be erroneous, but that does not affect the adequacy of the explanation. Since we *believe* that wishful thinking is an unreliable belief-forming process, we regard beliefs formed by wishful thinking as unjustified. What matters, then, is what we *believe* about wishful thinking, not what is *true* (in the long run) about wishful thinking. I am not sure how to express this point in the standard format of conceptual analysis, but it identifies an important point in understanding our theory.

III

Let us return, however, to the standard format of conceptual analysis, and let us consider a new objection that will require some revisions in the theory advanced until now. According to our theory, a belief is justified in case it is caused by a process that is in fact reliable, or by one we generally believe to be reliable. But suppose that although one of *S*'s beliefs satisfies this condition, *S* has no reason to believe that it does. Worse yet, suppose *S* has reason to believe that his belief is caused by an *un*reliable process (although *in fact* its causal ancestry is fully reliable). Wouldn't we deny in such circumstances that *S*'s belief is justified? This seems to show that our analysis, as presently formulated, is mistaken.

Suppose that Jones is told on fully reliable authority that a certain class of his memory beliefs are almost all mistaken. His parents fabricate a wholly false story that Jones suffered from amnesia when he was seven but later developed *pseudo*-memories of that period. Though Jones listens to what his parents say and has excellent reason to trust them, he persists in believing the ostensible memories from his seven-year-old past. Are these memory beliefs justified? Intuitively, they are not justified. But since these beliefs result from genuine memory and original perceptions, which are adequately reliable processes, our theory says that these beliefs are justified.

Can the theory be revised to meet this difficulty? One natural suggestion is that the actual reliability of a belief 's ancestry is not enough for justifiedness; in addition, the cognizer must be *justified in believing* that the ancestry

of his belief is reliable. Thus one might think of replacing (6_A), for example, with (7). (For simplicity, I neglect some of the details of the earlier analysis.)

(7) If *S*'s belief in *p* at *t* is caused by a reliable cognitive process, and *S* justifiably believes at *t* that his *p*-belief is so caused, then *S*'s belief in *p* at *t* is justified.

It is evident, however, that (7) will not do as a base clause, for it contains the epistemic term 'justifiably' in its antecedent.

A slightly weaker revision, without this problematic feature, might next be suggested, viz.,

(8) If *S*'s belief in *p* at *t* is caused by a reliable cognitive process, and *S* believes at *t* that his *p*-belief is so caused, then *S*'s belief in *p* at *t* is justified.

But this won't do the job. Suppose that Jones believes that his memory beliefs are reliably caused despite all the (trustworthy) contrary testimony of his parents. Principle (8) would be satisfied, yet we wouldn't say that these beliefs are justified.

Next, we might try (9), which is stronger than (8) and, unlike (7), formally admissible as a base clause.

(9) If *S*'s belief in *p* at *t* is caused by a reliable cognitive process, and *S* believes at *t* that his *p*-belief is so caused, and this meta-belief is caused by a reliable cognitive process, than *S*'s belief in *p* at *t* is justified.

A first objection to (9) is that it wrongly precludes unreflective creatures – creatures like animals or young children, who have no beliefs about the genesis of their beliefs – from having justified beliefs. If one shares my view that justified belief is, at least roughly, *well-formed* belief, surely animals and young children can have justified beliefs.

A second problem with (9) concerns its underlying rationale. Since (9) is proposed as a substitute for (6_A), it is implied that the reliability of a belief 's own cognitive ancestry does not make it justified. But, the suggestion seems to be, the reliability of a *meta-belief*'s ancestry confers justifiedness on the first-order belief. Why should that be so? Perhaps one is attracted by the idea of a 'trickle-down' effect: if an n+1-level belief is justified, its justification trickles down to an n-level belief. But even if the trickle-down theory is correct, it doesn't help here. There is no assurance from the satisfaction of (9)'s antecedent that the meta-belief itself is *justified*.

To obtain a better revision of our theory, let us re-examine the Jones case. Jones has strong evidence against certain propositions concerning his past. He doesn't *use* this evidence, but if he *were* to use it properly, he would stop

believing these propositions. Now the proper use of evidence would be an instance of a (conditionally) reliable process. So what we can say about Jones is that he *fails* to use a certain (conditionally) reliable process that he could and should have used. Admittedly, had he used this process, he would have 'worsened' his doxastic states: he would have replaced some true beliefs with suspension of judgment. Still, he couldn't have known this in the case in question. So, he failed to do something which, epistemically, he should have done. This diagnosis suggests a fundamental change in our theory. The justificational status of a belief is not only a function of the cognitive processes *actually* employed in producing it; it is also a function of processes that could and should be employed.

With these points in mind, we may tentatively propose the following revision of our theory, where we again focus on a base-clause principle but omit certain details in the interest of clarity.

(10) If S's belief in p at t results from a reliable cognitive process, and there is no reliable or conditionally reliable process available to S which, had it been used by S in addition to the process actually used, would have resulted in S's not believing p at t, then S's belief in p at t is justified.

There are several problems with this proposal. First, there is a technical problem. One cannot use an additional belief-forming (or doxastic-state-forming) process *as well as* the original process if the additional one would result in a different doxastic state. One wouldn't be using the original process at all. So we need a slightly different formulation of the relevant counterfactual. Since the basic idea is reasonably clear, however, I won't try to improve on the formulation here. A second problem concerns the notion of '*available*' belief-forming (or doxastic-state-forming) processes. What is it for a process to be 'available' to a cognizer? Were scientific procedures 'available' to people who lived in pre-scientific ages? Furthermore, it seems implausible to say that all 'available' processes ought to be used, at least if we include such processes as gathering *new* evidence. Surely a belief can sometimes be justified even if additional evidence-gathering would yield a different doxastic attitude. What I think we should have in mind here are such additional processes as calling previously acquired evidence to mind, assessing the implications of that evidence, etc. This is admittedly somewhat vague, but here again our ordinary notion of justifiedness is vague, so it is appropriate for our analysans to display the same sort of vagueness.

This completes the sketch of my account of justified belief. Before concluding, however, it is essential to point out that there is an important use

of 'justified' which is not captured by this account but can be captured by a closely related one.

There is a use of 'justified' in which it is not implied or presupposed that there is a *belief* that is justified. For example, if S is trying to decide whether to believe p and asks our advice, we may tell him that he is 'justified' in believing it. We do not thereby imply that he *has* a justified *belief*, since we know he is still suspending judgement. What we mean, roughly, is that he *would* or *could* be justified if he were to believe p. The justificational status we ascribe here cannot be a function of the causes of S's believing p, for there is no belief by S in p. Thus, the account of justifiedness we have given thus far cannot explicate *this* use of 'justified'. (It doesn't follow that this use of 'justified' has no connection with causal ancestries. Its proper use may depend on the causal ancestry of the cognizer's cognitive state, though not on the causal ancestry of his believing p.)

Let us distinguish two uses of 'justified': an *ex post* use and an *ex ante* use. The *ex post* use occurs when there exists a belief, and we say *of that belief* that it is (or isn't) justified. The *ex ante* use occurs when no such belief exists, or when we wish to ignore the question of whether such a belief exists. Here we say of the *person*, independent of his doxastic state vis-à-vis p, that p is (or isn't) suitable for him to believe.[17]

Since we have given an account of *ex post* justifiedness, it will suffice if we can analyze *ex ante* justifiedness in terms of it. Such an analysis, I believe, is ready at hand. S is *ex ante* justified in believing p at t just in case his total cognitive state at t is such that from that state he could come to believe p in such a way that this belief would be *ex post* justified. More precisely, he is *ex ante* justified in believing p at t just in case a reliable belief-forming operation is available to him such that the application of that operation to his total cognitive state at t would result, more or less immediately, in his believing p and this belief would be *ex post* justified. Stated formally, we have the following:

(11) Person S is *ex ante* justified in believing p at t if and only if there is a reliable belief-forming operation available to S which is such that if S applied that operation to his total cognitive state at t, S would believe p at t-plus-delta (for a suitably small delta) and that belief would be *ex post* justified.

For the analysans of (11) to be satisfied, the total cognitive state at t must have a suitable causal ancestry. Hence, (11) is implicitly an Historical account of *ex ante* justifiedness.

As indicated, the bulk of this paper was addressed to *ex post* justifiedness

This is the appropriate analysandum if one is interested in the connection between justifiedness and knowledge, since what is crucial to whether a person *knows* a proposition is whether he has an actual *belief* in the proposition that is justified. However, since many epistemologists are interested in *ex ante* justifiedness, it is proper for a general theory of justification to try to provide an account of that concept as well. Our theory does this quite naturally, for the account of *ex ante* justifiedness falls out directly from our account of *ex post* justifiedness.[18].

NOTES

[1] 'A Causal Theory of Knowing,' *The Journal of Philosophy* 64, 12 (June 22, 1967): 357–372; 'Innate Knowledge,' in S. P. Stich, ed., *Innate Ideas* (Berkeley: University of California Press, 1975); and 'Discrimination and Perceptual Knowledge,' *The Journal of Philosophy* 73, 20 (November 18, 1976), 771–791.

[2] Notice that the choice of a recursive format does not prejudice the case for or against any particular theory. A recursive format is perfectly general. Specifically, an explicit set of necessary and sufficient conditions is just a special case of a recursive format, i.e. one in which there is no recursive clause.

[3] Many of the attempts I shall consider are suggested by material in William P. Alston, 'Varieties of Privileged Access,' *American Philosophical Quarterly* 8 (1971), 223–241.

[4] Such a definition (though without the modal term) is given, for example, by W. V. Quine and J. S. Ullian in *The Web of Belief* (New York: Random House, 1970), p. 21. Statements are said to be self-evident just in case "to understand them is to believe them".

[5] Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1977, p. 22.

[6] I assume, of course, that 'nomologically necessary' is *de re* with respect to '*S*' and '*t*' in this construction. I shall not focus on problems that may arise in this regard, since my primary concerns are with different issues.

[7] This assumption violates the thesis that Davidson calls 'The Anomalism of the Mental'. Cf. 'Mental Events,' in L. Foster and J. W. Swanson, eds., *Experience and Theory* (Amherst: University of Massachusetts Press, 1970). But it is unclear that this thesis is a necessary truth. Thus, it seems fair to assume its falsity in order to produce a counter-example. The example neither entails nor precludes the mental–physical identity theory.

[8] Keith Lehrer's example of the gypsy lawyer is intended to show the inappropriateness of a causal requirement. (See *Knowledge*, Oxford: University Press, 1974, pp. 124–125.) But I find this example unconvincing. To the extent that I clearly imagine that the lawyer fixes his belief solely as a result of the cards, it seems intuitively wrong to say that he *knows* – or has a *justified belief* – that his client is innocent.

[9] This definition is not exactly what we need for the purposes at hand. As Ernest Sosa points out, introspection will turn out to be a belief-dependent process since sometimes the input into the process will be a belief (when the introspected content is a belief). Intuitively, however, introspection is not the sort of process which may be merely conditionally reliable. I do not know how to refine the definition so as to avoid this difficulty, but it is a small and isolated point.

[10] It may be objected that principles (6_A) and (6_B) are jointly open to analogues of the lottery paradox. A series of processes composed of reliable but less-than-perfectly-reliable processes may be extremely unreliable. Yet applications of (6_A) and (6_B) would confer justifiedness on a belief that is caused by such a series. In reply to this objection, we might simply indicate that the theory is intended to capture our ordinary notion of justifiedness, and this ordinary notion has been formed without recognition of this kind of problem. The theory is not wrong *as* a theory of the ordinary (naive) conception of justifiedness. On the other hand, if we want a theory to do more than capture the ordinary conception of justifiedness, it might be possible to strengthen the principles to avoid lottery-paradox analogues.

[11] I am indebted to Mark Pastin for this point.

[12] Cf. *Theory of Knowledge*, Second Edition, pp. 17, 114–116.

[13] Cf. 'A Causal Theory of Knowing,' *op. cit.* The reliability aspect of my theory also has its precursors in earlier papers of mine on knowing: 'Innate Knowledge,' *op. cit.* and 'Discrimination and Perceptual Knowledge,' *op. cit.*

[14] The view that introspection is retrospection was taken by Ryle, and before him (as Charles Hartshorne points out to me) by Hobbes, Whitehead, and possibly Husserl.

[15] Of course, if people in world *W* learn *inductively* that wishful thinking is reliable, and regularly base their beliefs on this inductive inference, it is quite unproblematic and straightforward that their beliefs are justified. The only interesting case is where their beliefs are formed *purely* by wishful thinking, without using inductive inference. The suggestion contemplated in this paragraph of the text is that, in the world imagined, even pure wishful thinking would confer justifiedness.

[16] I am indebted here to Mark Kaplan.

[17] The distinction between *ex post* and *ex ante* justifiedness is similar to Roderick Firth's distinction between *doxastic* and *propositional* warrant. See his 'Are Epistemic Concepts Reducible to Ethical Concepts?', in Alvin I. Goldman and Jaegwon Kim, eds., *Values and Morals, Essays in Honor of William Frankena, Charles Stevenson, and Richard Brandt* (Dordrecht: D. Reidel, 1978).

[18] Research on this paper was begun while the author was a fellow of the John Simon Guggenheim Memorial Foundation and of the Center for Advanced Study in the Behavioral Sciences. I am grateful for their support. I have received helpful comments and criticism from Holly S. Goldman, Mark Kaplan, Fred Schmitt, Stephen P. Stich, and many others at several universities where earlier drafts of the paper were read.