

## JUSTIFIED BELIEF AND EPISTEMICALLY RESPONSIBLE ACTION

Hilary Kornblith

One of the central projects of epistemological theorizing is the task of discovering the rules of ideal reasoning, both deductive and nondeductive. Another project involves appraising the beliefs of particular agents, or of particular communities. One obvious way to carry out this second project, assuming that the first has been completed, is to compare the manner in which the subject under study arrived at the belief in question with the rules of ideal reasoning. If the inferences made by the subject are licensed by the rules, then the belief is justified; if they are not, the belief is unjustified. Here we take a God's-eye view and hold our subject up to the highest standard of good reasoning.

Setting our standards this high is often precisely what we wish to do. It is often useful and interesting to see the extent to which a subject's patterns of reasoning differ from the ideal. In this kind of evaluation, we remove our subject from his historical context and compare his reasoning with that of an ideal subject, out of time, whose reasoning is uninfluenced by the fads and fashions of a community, and informed only by objectively right principles of reason. By appraising subjects' beliefs in this way, we are able to see the extent to which individuals or communities progress toward the ideal.

This is not, of course, the only standard by which one might reasonably appraise a subject's beliefs. Many subjects will, of course, fall short of the ideal. More importantly, subjects will often fall short of the ideal through no fault of their own. Rules of ideal reasoning are not easily come by, and it does not always show some shortcoming on the part of the subject that his reasoning was less than ideal. Sometimes we wish to know whether a subject was reasoning "as best he could," where this does not simply mean "in accord with rules of ideal reasoning"; we want to know whether the extent to which the subject departed from the ideal was his own fault.

What is at stake in this second kind of evaluation is a certain kind of responsibility: I will call it *epistemic responsibility*.<sup>1</sup> An *epistemically responsible agent* desires to have true beliefs, and thus desires to have his beliefs produced by processes which lead to true beliefs; his actions are guided by these desires. Sometimes when we ask whether an agent's belief is *justified* what we mean to ask is whether the belief is the product of *epistemically responsible action*, i.e. the product of action an epistemically responsible agent might have taken. I will use the term 'justified' throughout this paper in just this sense. My goal will be to further explicate this notion of justified belief and to explain the respects in which it differs from belief produced by inferences licensed by rules of ideal reasoning.

The motivation for investigating such a concept of justification is as follows. When we ask whether an agent's beliefs are justified we are asking whether he has done all he should to bring it about that he have true beliefs. The notion of justification is thus essentially tied to that of action, and equally to the notion of responsibility. Questions of justification are thus questions about the ethics of belief. In thinking about epistemology in ethical terms, however, it is important to remember that beliefs are not freely chosen; coming to believe something is not a voluntary action. The ethics of belief will thus not issue in rules of acceptance, but rather in rules of conduct. Truth-seeking agents ought to comport themselves in a certain manner. It is this perspective on justification that I will explore.

## I

The notion of justified belief, in the technical sense in which I am using this term, requires both more and less than belief produced by inferences licensed by rules of ideal reasoning. Let me begin by explaining the extent to which more is required.

---

<sup>1</sup>Laurence Bonjour introduces the term 'epistemic responsibility' in his "Externalist Theories of Justification," *Midwest Studies*, V (1980), 53–73. Bonjour, however, seems to presuppose that there is free choice of belief, and thus that fulfilling one's epistemic responsibility is a matter of following certain rules of ideal reasoning; I reject both of these assumptions. Bonjour is criticized for his assumption that there is free choice of belief by John Heil in his "Foundationalism and Epistemic Rationality," *Philosophical Studies*, forthcoming.

The manner in which one goes about acquiring evidence bears on the justificatory status of one's beliefs. An agent who reasons perfectly about evidence casually acquired, or who simply refuses to look at evidence not already in his possession, is acting in an epistemically irresponsible manner. Such beliefs may be perfectly supported by the evidence the agent has, yet because of the way in which the agent arrived at his belief, it seems unreasonable to say that he is justified in that belief. Thus, that appropriate logical relations among beliefs are psychologically instantiated in an agent does not show that the agent justifiably holds the resulting belief. The psychological instantiation of any theory of ideal reasoning is insufficient for justification.

My point here is simply this. Being justified requires more than simply reasoning properly; it requires that one gather evidence properly as well. In giving accounts of justification, epistemologists have often focused attention on proper reasoning to the exclusion of considerations about how evidence is gathered.<sup>2</sup> My objection to a theory of ideal reasoning as a theory of justification, in this section, is not that justified belief requires some kind of reasoning other than the ideal, but rather that justified belief requires more than just a constraint on reasoning.

It may be thought that cases of irresponsibly gathered evidence can be shown to be accommodated within the bounds of theories of ideal reasoning. These cases seem to fall into two categories: those in which the agent knows that the evidence he has was not gathered in a way which justifies his subsequent judgment, and those in which the agent is ignorant of this fact. Cases of the former kind are easily handled by ideal reasoning accounts, for this is merely another instance of the well-known phenomenon in which evidence  $e$  justifies belief that  $p$ , while  $e+e'$  does not. On the other hand, however, if the agent is ignorant of the defects in his evidence-gathering procedure, then the agent's conduct was not irresponsible and so his belief, as theories of ideal reasoning require, is justified.

I do not believe this reply to be adequate, for it presupposes that in epistemic contexts, unlike moral ones, ignorance always provides

---

<sup>2</sup>The tendency to hold an agent's evidence fixed and ask whether the agent's beliefs are justified relative to that evidence is a manifestation of this oversight.

one with an adequate excuse. On the contrary, I will argue that just as there are cases of morally culpable ignorance, there are cases of epistemically culpable ignorance as well.

Consider the case of Jones. Jones is a headstrong young physicist, eager to hear the praise of his colleagues. After Jones reads a paper, a senior colleague presents an objection. Expecting praise and unable to tolerate criticism, Jones pays no attention to the objection; while the criticism is devastating, it fails to make any impact on Jones' beliefs because Jones has not even heard it. Jones' conduct is epistemically irresponsible; had Jones' actions been guided by a desire to have true beliefs, he would have listened carefully to the objection. Since his continuing to believe the doctrines presented in his paper is due, in part, to this epistemically irresponsible act, his continued belief is unjustified. As we might expect from the parallel with moral evaluation of actions, this attribution of an unjustified belief reflects ill on Jones' character, or at last that part of his character which plays a role in the retention of this belief. If this aspect of Jones' character continues to play a significant part in his actions, more and more of his beliefs will be unjustified, in spite of the fact that the reasoning which led him to these beliefs, and for which he retains them, remains impeccable.<sup>3</sup>

Jones' belief is unjustified, after his colleague presents his objection, and it is unjustified because of his culpable ignorance. Since theories of ideal reasoning cannot account for the possibility of epistemically culpable ignorance, the requirements they specify are insufficient for justification. Justification, in the sense in which I am using the term, is not merely reasoning in accordance with principles of ideal reasoning.<sup>4</sup>

---

<sup>3</sup>This suggests a certain approach to cases dealing with evidence one does not possess. Gilbert Harman (*Thought* [Princeton: Princeton University Press, 1974], 142-154; and "Reasoning and Evidence One Does Not Possess," in P. French et al., eds., *Midwest Studies*, V [1980], 163-182) and William Lycan ("Evidence One Does Not Possess," *Australasian Journal of Philosophy*, 55 [1977], 114-126) have attempted to explain when evidence one does not possess defeats a knowledge claim. Although I am dealing with questions of justification, some of Lycan's discussion suggests an approach like mine. He speaks at one point (125), for example, of their being "epistemic obligations" to seek or to have certain kinds of evidence.

<sup>4</sup>While this point is universally noted in work on the methodology of science, it is rarely taken account of in epistemological work outside the

This point, it seems, can be easily handled. A theory of ideal reasoning needs to be supplemented by a theory of ideal evidence gathering in order to provide an adequate account of justification. This is, however, only one defensible *kind* of account of justification. Just as we can ask questions about justified belief *qua* ideally arrived at belief, we can ask questions about justified belief *qua* responsibly arrived at belief. In what follows, I will develop an account of justified belief as the product of epistemically responsible action, and I will show how it differs from belief arrived at through ideal reasoning and evidence gathering.<sup>5</sup>

## II

In determining whether a belief is justified, we must look at the process responsible for the presence of that belief. This involves, of course, looking at states of the believer, not merely at the instance at which we wish to know whether his belief is justified, but over some period of time. I believe that the period of time we must examine is much longer than is typically suggested. D.M. Armstrong, for example, has suggested that there is an important parallel between perceivers and thermometers.<sup>6</sup> Just as thermometers reliably register information about their environment in the form of pointers or columns of mercury, human perceivers (typically) reliably register information about their environment in the form of beliefs. This suggests, first, that if we are to investigate the process responsible for the presence of a perceptual belief we need look no further than the interaction of the perceiver with the object perceived, and, second, that the relevant process is one which happens to the epistemic agent, not one which involves any action on the agent's part.<sup>7</sup> Both of these suggestions are, I believe, false.

---

philosophy of science. Thus, while those working on analyses of knowledge and justification have, in recent years, tried to come to terms with the social character of knowledge, there is little discussion of the social character of justification.

<sup>5</sup>In what follows, I will use 'theories of ideal reasoning' as shorthand for 'theories of ideal reasoning and evidence gathering.'

<sup>6</sup>*Belief, Truth and Knowledge* (Cambridge: Cambridge University Press 1973), 166-169 and *passim*.

<sup>7</sup>These two suggestions are not discussed by Armstrong, nor is it clear that he is committed to accepting either of them. Thus, my point here is only that one should not be misled by the analogy.

Although it is clearly true that beliefs are not freely chosen, the actions which an agent freely chooses to perform may well affect the processes by which his beliefs are arrived at, and thus his beliefs themselves. It is thus that we may assess an agent, or an agent's character, by examining the processes responsible for the presence of his beliefs just as we may evaluate an agent, or his character, by examining the etiology of his actions. Actions which are the product of malice display a morally bad character; beliefs which are the product of epistemically irresponsible action display an epistemically bad character.

Even in the case of simple perceptual beliefs, which are, admittedly, for the most part arrived at automatically, the agent's actions play a role in the fine tuning of the belief acquisition process. An epistemically responsible agent must be on the uptake for defects in the process, and act to correct for them. While outright hallucination is rare, there is a wide range of familiar perceptual illusions which the epistemically responsible agent comes to take account of. A certain amount of accommodation to misleading experience is itself automatic and not the product of free action. Nevertheless, there can be no doubt that one can self-consciously instill in oneself a certain circumspection in circumstances where a mistake is likely to occur, such as in emotionally charged situations. Once this circumspection is acquired, more reliable belief acquisition will occur in one automatically, and without any particular action on the agent's part. Nevertheless, the presence of such automatic processes may often be traced to free action designed precisely to result in such processes, and when this is the case, it is to the agent's credit; failure to take such action may be epistemically irresponsible.

In determining the justificatory status of a belief, we must thus look beyond the process which gave rise to it and look at the means by which the process itself was arrived at. This will involve, for example, in perceptual cases, events long prior to the agent's interaction with the object perceived. This point allows us to appreciate a point about the logical form of justification statements.

### III

It is universally recognized that justification statements are time relative. An agent may be justified in believing that  $p$  at time  $t_1$ , and

yet unjustified in believing that  $p$  at time  $t_2$ . If the account of justified belief I offer is correct, however, justification statements are doubly time relative. Earlier I suggested that justified belief is belief which is the product of epistemically responsible action. Here I will make that account more precise: an agent is justified in his belief that  $p$  at time  $t$  as from time  $t'$  (where  $t'$  is earlier than  $t$ ) just in case all of the agent's actions between  $t'$  and  $t$  which affected the process responsible for the presence of the belief that  $p$  at  $t$  were epistemically responsible.

In examining the motivation for this view, it will be useful to consider an example. Mary is a logic student who believes that affirming the consequent is a truth-preserving argument strategy. In working out a logic problem, Mary carefully and judiciously applies those argument strategies she believes to be truth-preserving, including affirming the consequent. What is the justificatory status of the resultant belief?

Some distinctions are clearly called for here, and John Pollock has suggested that we may correctly describe the justificatory status of Mary's belief by distinguishing between an objective and a subjective sense of 'justification.'<sup>8</sup> Mary is objectively justified if her belief is arrived at on the basis of what are, in fact, good reasons; she is subjectively justified if her belief is arrived at on the basis of what she believes to be good reasons. Thus, on Pollock's account, Mary is not objectively justified, for her belief was not reached on the basis of principles of ideal reasoning; she is, however, subjectively justified, for her beliefs were arrived at on the basis of what, by Mary's lights, were good reasons.

Now there is no disputing what Pollock has to say about objective justification; Mary is not objectively justified. On the subjective side, however, I believe that Pollock is conflating a number of different cases. Consider Mary's belief that affirming the consequent is a truth-preserving argument strategy. She might have arrived at it in either of the following ways.

- (1) Mary's logic instructor, whom she had every reason to believe was at least minimally competent, told her that affirming the consequent is truth-preserving.

---

<sup>8</sup>"A Plethora of Epistemological Theories," in *Justification and Knowledge*, George Pappas, ed. (Boston: Reidel, 1979), 109/110.

- (2) On casual inspection, affirming the consequent seemed to Mary to be truth-preserving. Mary regularly applies rules which, on casual inspection, seem to her to be truth-preserving.

On Pollock's account, these two cases are on a par with each other; Mary is subjectively justified in her resultant belief in both cases. Nevertheless, there are important distinctions between these cases, and these distinctions can be made clear by recognizing that (subjective) justification statements are doubly time relative.

In case 2, for example, Mary's epistemic error lies in her casual acceptance of affirming the consequent as truth-preserving at  $t_0$ , not in her application at  $t_1$  of a rule she believes to be truth-preserving. Indeed, it would have been irresponsible of her, at  $t_1$ , having already made her mistake at  $t_0$ , not to apply the rule she believed to be truth-preserving. As from  $t_1$ , Mary's resultant belief at  $t_2$  (after having applied the rule) is justified, while as from  $t_0$ , the belief at  $t_2$  is not justified.

My treatment of these cases is exactly parallel to standard treatments of moral responsibility. In asking whether an action was morally justified, there are a number of different kinds of assessment which might be appropriate. We might wish to know whether the act in question was the sort which would have been performed by an ideal moral agent placed in the agent in question's situation. In assessing the degree to which an agent acted responsibly, however, this may not be the appropriate question to ask. When an agent acts judiciously on the basis of false moral beliefs, the same doubly time-relative assessment I favor in epistemic contexts is clearly called for.

Perhaps the kind of case which most clearly demands this treatment occurs when an otherwise morally depraved agent recognizes his depravity and decides to follow the path of the good and the right. Let us imagine a case in which the agent does not backslide, but, from the time of his decision, consistently performs actions which are morally good, as best he can determine. Given the state of his depravity at the time of his decision, not all of his subsequent actions are, objectively speaking, morally correct; yet we do the agent a disservice if we fail to recognize that, as from the time of his decision, all his actions are morally justified. Proper description of this example requires a doubly time-relative statement.



Epistemic contexts require the same treatment, and for much the same reason. Appropriate evaluation of an agent's epistemic conduct requires that when an agent does the best he can, epistemically speaking, in light of some previous error, this conduct be seen as on a par with neither conduct which is epistemically impeccable, nor conduct which is irresponsible through and through. This requires that we recognize the canonical form of justification statements to have two time indices.

## IV

An agent's belief at  $t_1$  may be arrived at by means which fail to accord with the principles of ideal reasoning as a product of epistemically irresponsible action. In such a case, there will be a time  $t_0$  prior to  $t_1$  such that the agent's belief at  $t_1$  is unjustified as from  $t_0$ . This may suggest the following account of the connection between justified belief and belief reached in accord with principles of ideal reasoning: if the belief that  $p$  at  $t_1$  is justified (i.e., responsibly arrived at) as from all  $t$ , the belief was arrived at in accord with principles of ideal reasoning.<sup>9</sup> This account is not correct.

It will not be sufficient to refute this thesis to show that there are patterns of reasoning, pervasive among adults, which are less than ideal, for epistemic irresponsibility may be rampant. Just as the pervasiveness of immoral conduct would be no evidence that one can act in a morally responsible way and yet perform immoral acts, the pervasiveness of less than ideal reasoning is not, by itself, evidence that such reasoning is arrived at in an epistemically responsible manner. In addition to showing that such reasoning is pervasive, I will show that it is arrived at responsibly. Perhaps the best example of this is inductive inference.

In making inductive inferences, people typically draw conclusions about a population based on a very small sample of that population. As Nisbett and Ross note in their *Human Inference: Strategies and Shortcomings of Social Judgment*, "People have little understanding of the relationship between the size of a sample and its faithfulness in reflecting the characteristics of the population from

---

<sup>9</sup>The converse of this account was shown to be false in section I above.

which it was drawn.”<sup>10</sup> Judgments about the characteristics of a population are frequently and confidently made on the basis of extraordinarily little information; in many cases, exposure to a single case is sufficient to result in a confidently held belief about a population. These are judgments which no reasonable normative theory could allow. In spite of the pervasiveness of this phenomenon, it is tempting to suppose that it is due to epistemic irresponsibility on a very large scale; even if there is some natural tendency to draw inductive conclusions on the basis of exposure to a small sample, this tendency would be checked in responsible adults. While I believe that the charge of pervasive epistemic irresponsibility can be supported in some cases, this is not such a case.

In most cases in which judgments are made about a population on the basis of a small sample, there is little or no opportunity to check one’s judgments about the population directly, that is, by checking the entire population. Indeed, if each of our judgments about a population were followed by exposure to information about the entire population, there can be little doubt but that our inferential habits would substantially improve. Given that such information is rarely easily available, we are forced to check our judgments against whatever piecemeal information we may come across. Needless to say, such a procedure is not likely to force re-evaluation of our methods of arriving at judgments.

It is well worth noting, as Nisbett and Ross point out,<sup>11</sup> that even careful attention to this kind of feedback will not always indicate the striking differences which one might expect between normatively correct and normatively incorrect decisions. Suppose you and I must each predict whether the next ball drawn from an urn will be black or white. We are each allowed to look at as many balls from the urn in advance as we like, though neither of us knows how many the other has examined. You decide to examine every ball in the urn; you predict black just in case there are more black

---

<sup>10</sup>(Englewood Cliffs: Prentice-Hall, 1980), 256. Nisbett and Ross are here summarizing the pioneering work of Tversky and Kahneman. See especially, Amos Tversky and Daniel Kahneman, “Belief in the Law of Small Numbers,” *Psychological Bulletin*, 76 (1971), 105-110 and “Judgment Under Uncertainty: Heuristics and Biases,” *Science*, 185 (1974), 1124-1131.

<sup>11</sup>*Op. cit.*, 256-260. The example below is adapted from these pages.

balls than white. I sample a single ball. Being hasty to generalize, I jump to the conclusion that the urn is filled almost entirely with black balls if I choose black, almost entirely with white balls if I choose white; I make my prediction accordingly. Clearly your prediction is made on ideal grounds while my judgment about the population is far less than ideal. Nevertheless, if we play this game only once, I am not likely to do much worse than you.

Consider a simple case. Suppose the urn is filled with 10 per cent black balls and 90 per cent white ones. You will count them all out, and, predicting that the next ball drawn will be of the color of the majority, you predict that it will be white. You will be right, of course, 90 per cent of the time. I pull out a single ball and, because I believe that nearly all the balls are of that color, I predict that the next ball drawn will be of that color. My chances of being correct are  $((.90 \times .90) + (.10 \times .10)) = 82$  per cent. This is not a very large difference. The chances that you will be right and I wrong are  $(.90 \times .18) = 16.2$  per cent. The chances that I will be right and you wrong are  $(.82 \times .10) = 8.2$  per cent. It is quite clear that if we play the game only once when the ratio of black to white is 1 : 9, I am unlikely to see the error of my ways. Chances are we will both be right in our predictions. There is even some chance that I will be right and you wrong, and this will only convince me that my decision procedure is a good one.

Matters are even worse if the ratio of black balls to white balls is more nearly even. Consider the limiting case where the urn is half filled with black balls and half filled with white. You predict that the next ball to be drawn will be white (since you predict black just in case there are more black balls than white), and you will be correct 50 per cent of the time. My prediction will be correct  $((.5 \times .5) + (.5 \times .5)) = 50$  per cent of the time. You will be right and I wrong  $(.5 \times .5) = 25$  per cent of the time; I will be right and you wrong  $(.5 \times .5) = 25$  per cent of the time. I am unlikely to be convinced by this demonstration that your judgment about the population was more reasonable than mine. Intermediate cases do not show large differences either.

Of course, it is psychologically unrealistic to suppose that I would believe a single ball drawn from an urn to be representative of the entire population. Such jumping to conclusions is not at all uncommon, however, in less contrived situations. To cite a single exam-

ple, we often form judgments about a wide variety of consumer products on the basis of a single person's bad experience. What the story of the urn indicates is that the results of our radically less than ideal belief about a population on the basis of a single case may not make plain to us just how much less than ideal our belief was.

What this means is that agents who have a tendency to make inductive inferences on the basis of small samples, and this includes most everyone, are extremely unlikely to detect their errors. While the judgments typically made about a population based on a small sample are clearly not in accord with principles of ideal reasoning, whatever one's view of ideal reasoning may be, the procedures which we are typically in a position to apply in order to test these judgments are not ones which are likely to expose our errors. One can hardly fault an agent then for making judgments about a population on the basis of a small sample, when there is a natural tendency to make such judgments and little opportunity to detect one's errors. Here is a case then in which epistemically responsible conduct leads to less than ideal reasoning.

It should not be surprising that there are such cases. While epistemically responsible conduct can be expected to eliminate gross errors in reasoning, small deviations from the ideal may elude even the epistemically responsible agent. The connection between epistemically responsible conduct and ideal reasoning is thus not a simple one.

## V

It will also be worth comparing my view with the view that justified belief can be identified with reliably produced belief.<sup>12</sup> Although some may wish to identify reliably produced belief with ideal reasoning, and will thus take it as already established that justified belief cannot be identified with reliably produced belief, I wish to remain neutral on the relationship between ideal reasoning and reliability. Whatever their relationship, justified belief cannot be identified with reliably produced belief.

This point is demonstrated by an application of the example in the last section. Beliefs formed about a population on the basis of a

---

<sup>12</sup>As suggested by Alvin Goldman, "What Is Justified Belief?" in George Pappas, ed., *op. cit.*

small sample are responsibly arrived at, and thus justified, but they are not formed by a process which tends to produce true beliefs. Although reliably produced belief may be an ideal to be sought, it cannot be a requirement for justified belief. Since epistemically responsible action may result in something less than reliably produced belief, an agent may be justified in holding a belief without that belief being reliably produced. Beliefs produced by unreliable processes, where the extent of the unreliability would not be detected by an epistemically responsible agent, are nonetheless justified. Similarly, beliefs produced by reliable processes, where the existence of the process is due to epistemically irresponsible action, fail to be justified.

Similar considerations apply to Richard Boyd's suggestion that justified belief be identified with reliably regulated belief.<sup>13</sup> While I agree with Boyd that the notion of justification is more closely tied to the regulation of belief than to its production, the discussion in Section IV suggests that even the regulation of belief acquisition processes need not be reliable for justified belief. Our ability to regulate our acquisition of beliefs is both aided and limited by our beliefs and dispositions. Since even the epistemically responsible agent's beliefs need not be true and his dispositions need not be fully reliable, his resulting ability to regulate the processes by which his beliefs are acquired will be less than optimal. Again, reliably regulated belief is an ideal to be sought, but it cannot be a requirement for justified belief.

In spite of these objections, I think it is important to see that the account I offer is very much in the spirit of the naturalized epistemology of Goldman and Boyd, and, indeed, allows us to solve an important problem facing naturalistic accounts.<sup>14</sup> According to reliability theorists, a belief producing process need only be reliable in order for it to be justification conferring; no beliefs about the

<sup>13</sup>"Scientific Realism and Naturalistic Epistemology," *Proceedings of the Philosophy of Science Association*, 2 (1980).

<sup>14</sup>Indeed, the problem is really more general than this. As William Alston suggests ("Chisholm's Epistemology," *Noûs*, XIV [1980], 584), ". . . it is self-defeating to restrict justifiers to facts that are known by, or evident to, the believer. That restriction will wash out practically all the plausible candidates, not just those of some particular position." Thus, the question facing epistemologists is not 'Are there justifiers which need not be known by the believer?' but rather 'Which justifiers need not be known by the believer?'

reliability of the process are necessary. An example due to Laurence Bonjour, however, seems to show that there are cases in which a reliable process fails to be justification conferring just because it is not believed to be reliable.<sup>15</sup> As Bonjour points out, someone who is in fact clairvoyant, yet both fails to have any reason to believe that he is clairvoyant and does have reason to believe that he is not, is not justified in accepting the deliverances of his special sense. Similarly, *contra* Boyd, cases can be constructed in which an agent's beliefs are reliably regulated, though accidentally so, and thus unjustified. The problem seems to be to separate the reliable belief-forming processes, or reliable belief-regulating processes, which are justification-conferring, from those which are not. If my account is correct, a proper solution to this problem forces us to recognize the irrelevance of the actual reliability of either belief-producing or belief-regulating processes.

The story I have to tell about justification is a naturalistic one. We are born with certain belief-forming processes and dispositions instrumental in regulating these processes. These are both at least roughly instrumental in getting at the truth, but, on my account, this has nothing to do with justification itself. On my account, having justified beliefs is simply doing the best one can in light of the innate endowment one starts from, however reliable or unreliable it may be. This is not to say that the fact that our innate endowment is at least roughly reliable is irrelevant to epistemology, for if one wishes to explain why justified beliefs are likely to be true, one must appeal to the reliability of our innate endowment.

There are some belief-forming processes which, given the human psychological constitution and the standard range of human experience, could not be arrived at in an epistemically re-

---

<sup>15</sup>Bonjour, *op. cit.* This is not quite fair to Goldman, however, for his notion of relevant alternatives seems to provide a way of specifying which processes need be backed by beliefs about their reliability in order to be justification conferring. See Goldman, "Discrimination and Perceptual Knowledge," *Journal of Philosophy*, LXIII (1976), 771-791; Hilary Kornblith, "Beyond Foundationalism and the Coherence Theory," *Journal of Philosophy*, LXVII (1980), 597-612 and "The Psychological Turn," *Australasian Journal of Philosophy*, LX (1982). This does not, however, solve the problem; it only gives the form of a solution. The only way I now see of filling in this outline involves downplaying the importance of reliability significantly. See below.

sponsible manner unless they were believed to be reliable; for other processes this is not so. The standard here must be that of the epistemically responsible agent; there is no way to judge this issue *a priori*. If I am correct, the question of reliability thus drops out entirely.

VI

Some writers define justification in terms of truth. Alvin Goldman is the most radical here; he defines justified belief as belief which is reliably produced: belief which is produced by a process which tends to yield true beliefs. Other writers see no such close connection between justification and true belief.<sup>16</sup>

On my account, there is a connection between justification and truth, but it is not nearly so close as Goldman suggests. Justified belief is belief which is the product of epistemically responsible action; epistemically responsible action is action guided by a desire to have true beliefs. The epistemically responsible agent will thus desire to have true beliefs and thus desire to have his beliefs pro-

---

<sup>16</sup>I do not intend, here, to be drawing the distinction between internalist and externalist theories of justification, but it is worth asking how the view proposed maps on to that distinction. Unfortunately, there is no uncontroversial way to draw the distinction. I believe that the following is not unreasonable; it clearly shows at least why we should regard Descartes as an internalist and Goldman as an externalist. Internalists believe both that (1) there is a privileged type of epistemological starting point, the benefits of which over other possible types of starting points can be shown on internal grounds, i.e. in virtue of features of that starting point accessible to the agent; and (2) whether an agent is justified in holding a belief can always be determined by the agent in virtue of features of his position accessible to him. Externalists deny both these two points. On this way of drawing the distinction, the view proposed here is a hybrid. With externalists, I deny (1), for on my conception of justification, being justified is simply doing the best one can given one's innate endowment, whatever that may be. One's epistemological starting point thus has no benefits over other possible types of starting points, from the point of view of justification. On the other hand, I agree with internalists on (2), since whether one is acting in an epistemically responsible manner is accessible to the agent. For a more complete defense of the view that a proper theory of justification lies somewhere between internalism and externalism, see Andrew Altman and Michael Bradie, "A Contextualist Critique of Goldman's Epistemology," presented at the 1981 Western Division Meetings of the APA. No doubt there are other ways of drawing this distinction on which my view will turn out not to be a hybrid.

duced by reliable processes; but even doing the best he can, this desire need not be fulfilled. The further the agent's innate intellectual endowment is from being reliable, the further will his justified beliefs be from being reliably produced. We can at least conceive of creatures with wholly unreliable innate endowments, and, consequently, with massively false systems of justified belief.

We can be quite confident that we are not such creatures. As Quine states, "There is some encouragement in Darwin."<sup>17</sup> But while Darwin provides us with encouragement here, psychological research shows us the limits of our innate endowment. Our systems of belief acquisition are less than fully reliable, and may even be less than fully reliable in the long run. Even though the encouragement we take from Darwin suggests that epistemically responsible action will likely result in more nearly reliably produced belief, genuine reliability remains an ideal, and not the standard against which justified belief is to be measured.<sup>18</sup>

*University of Vermont*

---

<sup>17</sup>"Natural Kinds," in *Ontological Relativity and Other Essays* (New York: Columbia University Press, 1969), 126.

<sup>18</sup>My approach here draws heavily on two recent papers by Alvin Goldman: "Epistemics: The Regulative Theory of Cognition," *Journal of Philosophy*, LXXV (1978), 509-523; "Epistemology and the Psychology of Belief," *The Monist*, 61 (1978), 525-535. Ernest Sosa suggests a project very similar to the project of this paper in the final paragraphs of his "The Raft and the Pyramid," *Midwest Studies*, V (1980), 23. I have received helpful comments from Alvin Goldman, Patricia Kitcher, Philip Kitcher, Arthur Kuflik, David Shatz, Fred Schmitt, George Sher and the editors of the *Philosophical Review*. I have also profited from discussions on this topic with Richard Boyd. Work on this project was supported by a University of Vermont Summer Research Grant and the National Endowment for the Humanities.