teorema
Vol. XX/3 (2003) pp. 83-91

Truth and the Aim of Epistemic Justification*

Hamid Vahid

RESUMEN

Cualquier teoría de la justificación epistémica tiene que enfrentarse a la cuestión de cuál es su objetivo y por qué lo valoramos. El consenso casi general entre los epistemólogos es que existe una conexión íntima entre justificación y verdad. Sin embargo, la formulación estándar de esta conexión (como la meta dirigida hacia la verdad de maximizar la creencia verdadera y minimizar la falsa) ha sido desafiada en los últimos tiempos, de manera cada vez más insistente, sobre la base de que no deja lugar para creencias falsas justificadas y creencias verdaderas no justificadas. En este artículo voy a argumentar a favor de una versión diacrónica de la meta dirigida hacia la verdad a la vista de las objeciones que se han levantado en contra de ella. Para proporcionar un soporte adicional a esta conclusión, critico la explicación sincrónica de Foley de la meta dirigida hacia la verdad mostrando cómo resulta socavada por ciertas consecuencias de su teoría egocéntrica de la racionalidad epistémica.

ABSTRACT

Any theory of epistemic justification must address the question of what its aim is and why we value it. The almost general consensus among epistemologists is that there is an intimate link between justification and truth. However, the standard formulation of this connection (as the truth-directed goal of maximizing true belief and minimizing false belief) has been increasingly challenged in recent times on the ground that it leaves no room for justified false beliefs and unjustified true beliefs. In this paper, I shall argue for a diachronic version of the truth-directed goal in the face of the objections raised against it. To provide further support for this conclusion, I criticize Foley's synchronic account of the truth-directed goal by showing how it is undermined by certain consequences of his egocentric theory of epistemic rationality.

Beliefs can be evaluated from a number of perspectives. Depending on our choice of the standards and goals (moral, practical, etc.) the evaluation will yield different results. Epistemic evaluation, however, involves epistemic standards and appropriate epistemic goals. A theory of epistemic justification must, thus, address the question of the aim and objective of epistemic justification i.e., what is the point of epistemic justification and why we value it. It is generally thought that there is an intimate connection between justification and truth. This is usually construed along the lines that a belief is justified if forming that belief is a good thing from the epistemic point of view

which, in turn, is characterized in terms of the aim of maximizing truth and minimizing falsity in a large body of beliefs (call this, following Foley, the 'truth-directed goal'; See, for example, BonJour (1985), Alston (1989) and (1996) and Foley (1987)).

I. THE TRUTH-DIRECTED GOAL: RESPONDING TO OBJECTIONS

In this paper, I shall defend a diachronic version of the truth-directed goal in the face of the some of the objections raised against it. I shall begin with a prominent objection that was first fully stated in a paper by Maitzen and has since being endorsed by a number of epistemologists (including David [David (2001)]). The objection is quite simple; the standard formulation of the truth-directed goal leaves no room for justified false beliefs and unjustified true beliefs.

> If justification is essentially a matter of serving [the truth-directed goal], then it seems we would evaluate no true beliefs as unjustified and no false belief as justified [...]. Therefore, if […the truth-directed goal] governs epistemic justification, no instance of true belief will fail to be justified unless it tends *not* to maximize true belief and minimize false belief [...]. But how could S's true belief at t have anything but a verific *tendency*? [...]. In sum then, S's believing that p at t will be justified unless somehow it tends not to maximize, in the long run, her stock of true beliefs. But without much more argument, I can't see how S's belief at t of an arbitrary true proposition could have such a tendency [Maitzen (1995), p. 872].

Before turning to the above objection, it is necessary to be clear about one point. The truth-directed goal does not involve individual beliefs so as to make a justified belief come out always as true and an unjustified belief come out always as false. Rather the aim of justification is to have an accurate and comprehensive belief system. By focusing on belief systems rather than individual beliefs an initial hurdle on the way of obtaining false but justified beliefs and true but unjustified beliefs is removed. The situation is rather analogous to the use of the Davidsonian principle of charity in interpreting speakers of a language. For the principle of charity seems to make it impossible to interpret a speaker as ever having a false belief (by our lights). Once we go directly from facts about which sentences someone holds true under certain circumstances to conclusions about the interpretation of these sentences, the principle would seem to rule out the possibility of error. For this reason, Davidson changes the form of evidence that is needed for the theory of interpretation by focusing on languages and communal attitudes to the truth of sentences. Principle of charity is, thus, recognized as holistic constraint applying not to individual beliefs, but rather to *systems* of beliefs. We must interpret speaker so that most of the beliefs in his systems are true (by

our lights), leaving room for ascribing some false beliefs to him. Let us now return to Maitzen's argument.

To begin with we need to note that, unlike how Maitzen and others treat the problem, there are actually two distinct (though related) questions involved in the above objection: (1) How are justified false beliefs and unjustified true beliefs possible on a truth-directed goal of justification, and (2) if these are possible, how in the case of, say, justified false beliefs, believing a false proposition could maximize one's stock of true beliefs? I start with the first question. First a general remark. The aim of justification being to maximize true beliefs and minimize false beliefs is compatible with justification failing to reach that aim i. e., for a justified belief to turn out to be false and an unjustified belief to turn out to be true. This is rather analogous to the case of functions. Items may have functions which they do not perform. A heart in fibrillation and the eyes of a congenitally blind person, still have their function. Any theory of function must be able to explain how malfunction is possible. Likewise, any theory of justification must be able to explain how some justified beliefs turn out to be false and some unjustified beliefs turn out to be true.

I think both of these possibilities can be explained by reference to the perspectival character of epistemic justification. Unlike truth, justification is perspective i.e., it is determined relative to the cognizer's evidence. If a belief is true, then it is true at all times. Its truth-value remains constant across time or population. By contrast, the justification of a belief depends on the epistemic circumstances of the cognizer holding that belief. Depending on the body of evidence he possesses, the attribution of justified belief varies accordingly. Thus, an agent may be unjustified in holding a belief p at t_1 because he lacks the relevant evidence, whereas the very same belief might be justified for him or another person at t_2. The perspectival character of epistemic justification is one of its non-negotiable features, and it is precisely this characteristic that obviates the collapse of epistemic justification into truth [Alston (1996), ch. 8]. Justification is a means to truth, not identical with it. Now, given the truth-directed goal of justification, we wish to maximize our stock of true beliefs and minimize our stock of false beliefs, but, in view of the perspectival character of justification, we are *bound* to do it relative to the body of evidence, e, at our disposal. The epistemic situation of a cognizer is, however, often incomplete or defective i.e., it fails to include all the relevant evidence. So e might be equally incomplete or defective. Now two possibilities might arise vis-á-vis e and a belief p. Relative to e, the belief p may be more likely to be true and so, on the truth-directed goal, justified where it is in fact false. This results in a case in which a justified belief is false. On the other hand, it may that, relative to e, p is more likely to be false, and, thus, unjustified, whereas it is, in fact, true. This would be a situation in which an unjustified belief is true. Only under ideal circumstances i.e., circumstances in which all relevant evidence is available, can a belief be justified if and

only if it is true. But our epistemic predicament often falls short of being ideal. So, given the perspectival character of epistemic justification and its truth-directed goal, there would be room for unjustified true beliefs and justified false beliefs.

II. THE TRUTH-DIRECTED GOAL: SYNCHRONIC VS. DIACHRONIC VERSIONS

We can now turn to the second question we discerned in Maitzen's argument: If having an unjustified true belief is possible on a truth-directed goal of justification, then it must tend not to maximize, in the long run, one's stock of true beliefs, but how could the agent's "belief at *t* of an arbitrary true proposition [...] have such a *tendency*?" [Maitzen (1995), p. 872]. This question actually arises from a synchronic understanding of the truth-directed goal. If we regard it only as a function of truth-ratio in a body of belief at a time *t*, then the sense of bewilderment is justified. But the truth-directed goal is a historical concept. It is the overall maximization of truth and minimization of falsity in one's belief repertoire in the long run that it is concerned with. However, this overall effect comes by only if certain antecedent conditions of belief formation are satisfied. It is, thus, not just the truth-ratio that determines whether believing *p* is a good thing form the epistemic point of view but also how the belief is formed or sustained. We are, to repeat, concerned with the overall maximization of truth in our belief system not just a particular time slice of it. The explanation of how believing an arbitrary true proposition could fail to maximize our stock of true beliefs can now be stated as follows. Since justification is perspectival, the available evidence *e* at *t*, being incomplete, might indicate that our belief *p* is more likely to be false, even if, as it turns out, it is in fact true. Relative to *e*, the belief *p* is not adequately grounded. So believing *p* is not a good thing from the epistemic point of view, and thus, unjustified. However, we have treat like cases alike. This means that in the future we have to treat other inadequately grounded beliefs as unjustified. As a norm, however, most of the inadequately grounded beliefs turn out to be false. It is only by accident that the belief *p* turns out to be true. This is how believing a true proposition can fail to bring about an overall maximization of truth in our belief repertoire. Truth-directed goal is not just a matter of the truth-ratio in our belief system but also of how these beliefs are formed.

The same goes for the case of justified false beliefs. One might wonder how accepting a false proposition could tend to maximize truth in one's belief system. The explanation proceeds along the same line as in the previous case. Given the perspectival character of justification, our available evidence *e* might indicate that the belief *p* is more likely to be true, whereas it is in fact

false. Relative to *e* believing *p* is a good thing from the epistemic point of view because it is adequately grounded and, thus, justified. Again we have to treat like cases alike i.e., take the well-grounded beliefs in the future as justified as well. As a norm, however, most of the well-grounded beliefs turn out to be true. This is how believing a false proposition can result in the overall maximization of true beliefs in our belief stock. It is not just the truth value of a belief that counts towards its justification (vis-à-vis the truth-directed goal) but how the belief has been formed in a non-accidental and reliably way. The synchronic truth-ratio version of the truth-directed goal ignores the way a belief has been formed.

In responding to Maitzen, I have appealed to the claim that we have to treat like cases alike e.g., if we regard a belief as justified because it is adequately grounded, we have to regard other adequately grounded beliefs as justified as well. What is the rationale behind this claim? Well, the rationale involves the so-called thesis of epistemic supervenience according to which epistemic properties (e.g., being justified) supervene on non-epistemic properties such as being reliably formed, or being adequately grounded. If we have two beliefs which are adequately grounded and one is justified so is the other. So to say that believing *p* is a good thing from the epistemic point of view of maximizing true beliefs is just to say that believing *p* has a tendency to bring about an overall maximization of true beliefs in our belief repertoire. This tendency is ensured by the "fact" that justification is a supervenient concept requiring us to universalize our epistemic judgement about a particular case to other similar cases.

I think it is the supervenient character of epistemic justification that rules out the synchronic conception of the truth-directed goal which is concerned only with the truth-ratio of our belief system at a particular time slice. The truth-directed goal is, rather, of a diachronic nature as it involves the overall maximization of true beliefs and minimization of false beliefs in our belief repertoire that, in turn, includes the past as well as the future history of our belief-forming practices. Marian David has objected that such a diachronic account is troubling because "[i]t is very implausible to hold that whether I am now justified in believing *p* depends partly, or wholly, on the truth values of my future beliefs" [David (2001), p.166]. But this objection is irrelevant. From the fact that the truth-directed goal is of a diachronic nature, it does not follow that whether I am justified in believing *p* depends on the truth value of my future belief *q*. It is not the case that the truth value of one belief enters into the body of evidence that is exploited to evaluate the epistemic status of another. Rather, the epistemic standings of the two beliefs are connected to one another through the intervention of a third factor, namely, the way they are produced. Belief *p* is justified, we assume, in virtue of being produced by a reliable process. The same is true of belief *q* (if it is also reliably produced). This is in fact what the thesis of epistemic supervenience entails.

In a different way Michael DePaul has criticized the truth-directed goal on the ground that if we take truth as being our only epistemic goal, then we would not be able to explain how knowledge is better than mere true belief, "But if we are pluralists about epistemic value and take justification to be good apart from any connection it might have to the truth, then we can explain the superior value of knowledge" [DePaul (2001), p. 180]. To give an example, he says, it is very reasonable to value a victory obtained by exercising a brilliant strategy more than a victory brought about by luck. But this is because, he says, more than one value is at stake here. In addition to taking victory as a good thing, we also regard excellence and competence as good as well. Likewise for the case of valuing knowledge over mere true belief. But this value pluralism is plausible only if we focus on a particular time slice of our goals. So long as our aim is to win a battle at a time *t*, then preferring a well-conducted victory over a haphazard and accidental one requires we take two values into account i.e., victory and competence. But if our aim is "wining battles in the long run" not just at a particular time t, then, to prefer a well-conducted victory over an accidental one, we need not take up value pluralism. The aim of obtaining victory in the long run necessarily requires preferring a well-conducted victory. Likewise, we value knowledge because we want to have true belief in the long run rather than just at a particular moment. Once we adopt a diachronic conception of the aim of justification as overall maximization of truth (or "truth in the long run"), we can ensure that knowledge is more valuable than mere true belief without embracing value pluralism.

It might be objected (as both Maitzen and David have) that, given the preceding remarks, our epistemic aim is no longer believing true and avoiding false propositions, but, rather, having adequately grounded or reliably formed beliefs. But even if we regard the goal of justification as forming beliefs by way of reliable mechanisms, that still does not dispose of the truth-directed goal. For any feature of beliefs that is probably correlated with truth, counts towards the original aim of maximizing true and minimizing false beliefs. These factors include, reliable belief formation, being the best explanation of a certain phenomenon, etc. To deny this is, as we said, to understand the aim of maximizing true beliefs in much too crude a sense. As I have emphasized, by being sensitive to how a belief has been formed, the aim of justification would have a diachronic nature. One clearly historical account of justification is reliabilism which, as David points out, "has to maintain that the truth-[directed] goal is the diachronic goal of having true beliefs in the long run" [David (2001), p.166]. The reason is obvious. By taking the past as well as the future performance of belief-forming processes into account, reliabilism naturally promotes a diachronic understanding of the truth-directed goal. To reinforce the conclusions reached so far, I shall now proceed to consider Foley's account of rationality which explicitly advocates a synchronic

version of the truth-directed goal. I shall show that the 'reliabilist' ingredients of Foley's account undermine his synchronic reading of the truth goal.

Foley regards rationality (justification) as a goal-directed notion and defines it in terms of the point of view one adopts [See Foley (1987) and (1993)]. This point of view consists of a goal, a perspective and a set of resources. A perspective is essentially a set of beliefs, and what perspective one adopts in making a rationality judgement depends on one's interests as well as the context. Resources consist of information and the data represented in one's beliefs. These vary in accordance with theory of justification under question. Foley then articulates his radically subjective conception of epistemic rationality along the following lines. He takes as his goal, the synchronic goal of now believing those propositions that are true and now not believing those propositions that are false. The perspective is an egocentric one, namely, that of individual believers. So it is egocentrically rational for you to believe a proposition only if you would think on deep reflection that believing it is conducive to having an accurate and comprehensive belief system. Thus, Foley construes egocentric rationality as depending subjunctively on reflection in the sense that an agent need not have actually engaged in reflecting on a proposition in order to be justified in believing it. Rather, it is enough for his current belief to conform to what his stable opinion would be. He gives an example in which a man and a woman, who are both ideally reflective, come to believe a proposition on exactly the same evidence. However, he not she would come to believe that the evidence in question is inadequate (upon reflection). According to Foley this counterfactual difference affects their epistemic status regardless of whether this difference is reflected in their current beliefs. This difference is enough to make her, rather than him, egocentrically rational to believe that proposition.

As can be seen, reflection plays an important role in Foley's theory of rationality. However, not any old notion of reflection suits his scheme. To begin with, reflection should be deep. The agent has to go on till further reflection can no longer affect his judgement. This is where we reach the point of stability, the point when we become invulnerable to self-criticism. He further requires that during the process of reflection the agent should not be subject to distorting influences such as drowsiness and drunkenness. According to Foley, under such conditions reflection would be indicative of an agent's deepest epistemic standards. So the rationality of a belief depends on how it would be assessed in terms of the agent's deepest epistemic standards (which would be naturally construed in terms of the truth-directed goal).

However, having relativized rationality to an agent's point of view, the process of reflection, with the kind of constraints that Foley imposes on it, emerges as functioning like a reliable process. It will not obviously be a reliable indicator of the truth of a belief (as in a proper reliability theory), but, given Foley's account, a reliable indicator of an agent's deepest epistemic

standards (recall Foley's example of the man and the woman). So despite the obvious fact that Foley's account differs from a normal reliability theory of justification in being subjective, it shares with that theory reliance on the contribution of reliable processes to the justification of one's belief (though, as emphasized, the reliability of reflection, on his account, is indirectly linked to the truth of the belief by revealing the agent's deepest epistemic standards). It is also, as with the reliability theory, an externalist account of justification. Foley quite explicitly distinguishes between, what he calls, 'epistemic' and 'metaphysical' varieties of internalism, with the latter requiring that the conditions of rational belief be internal i.e., taking place from the skin in. He advocates only the metaphysical version of internalism adding that "it is not a version of internalism that implies you already have access to the conditions that make your belief rational or irrational" [Foley (1993), p. 101]. So he is really an externalist in the sense relevant to how the internalism/externalism controversy is usually perceived. Now, if for Foley, the process of reflection is to be reliable process, then, with reliability being a function of the process's past and future performance, the truth-directed goal cannot be, as he claims, of a synchronic nature. I conclude, therefore, that a diachronic version of the truth-directed goal has the best chance of linking the notions of epistemic justification and truth while avoiding the implausible consequences that are thought by some theorists to follow from the thesis of the truth goal.

Analytic Philosophy Faculty, IPM
P.O. Box 19395-1795,
Tehran, Iran.
E-mail: vahid@ipm.ir

NOTES

* I would like to thank Richard Foley for helpful comments on an earlier draft of this paper.

REFERENCES

ALSTON, W. (1989), *Epistemic Justification*, Ithaca, Cornell University Press.
— (1996), *A Realist Conception of Truth*, Ithaca, Cornell University Press.
BONJOUR, L. (1985), *The Structure of Empirical Knowledge*, Cambridge, Mass., Harvard University Press.
DAVID, M. (2001), "Truth as the Epistemic Goal", in Steup (2001).
DEPAUL, M. (2001), "Value Monism in Epistemology", in Steup (2001).
FOLEY, R (1987), *The Theory of Epistemic Rationality*, Cambridge, Mass., Harvard University Press.
— (1993), *Working Without a Net*, Oxford, Oxford University Press.

MAITZEN, S. (1995), "Our Errant Epistemic Aim", *Philosophy and Phenomenological Research* 55, pp. 869-76.
STARTWELL, C. (1992), "Why Knowledge is Merely True Belief", *Journal of Philosophy* 89, pp. 167-80.
STEUP, M. (2001), *Knowledge, Truth and Duty*, Oxford, Oxford University Press.