Jesús P. Zamora Bonilla

TRUTHLIKENESS WITH A HUMAN FACE

ON SOME CONNECTIONS BETWEEN THE THEORY OF VERISIMILITUDE AND THE SOCIOLOGY OF SCIENTIFIC KNOWLEDGE

ABSTRACT. Verisimilitude theorists (and many scientific realists) assume that science attempts to provide hypotheses with an increasing degree of closeness to the full truth; on the other hand, radical sociologists of science assert that flesh and bone scientists struggle to attain much more mundane goals (such as income, power, fame, and so on). This paper argues that both points of view can be made compatible, for (1) rational individuals only would be interested in engaging in a strong competition (such as that described by radical sociologists) if they knew in advance the rules under which their outcomes are to be assessed, and (2), if these rules have to be chosen "under a veil of ignorance" (i.e., before knowing what specific theory each scientist is going to devise), then rules favoring highly verisimilar theories can be prefered by researchers to other methodological rules.

The theory of verisimilitude is a theory about the aim of science. In a well known paper (Popper 1972), written before developing his own approach to the topic of verisimilitude, Popper described that aim as the production of testable explanations of whatever facts we thought to be interesting to explain, though he also recognised that it was rather improper to talk about the aims *of science*, since only *scientists* have goals, properly speaking, and these may look for a wide variety of things. Most discussions about the concept of truthlikeness have obviously been concerned with the first of these questions – say, what is the cognitive goal of science, assuming that one such goal exists – but they have largely ignored the second one, i.e., what the connection may be between that epistemic goal and the actual motivations and behavior of scientists. In this brief paper I would like to make a contribution to the second topic, though the ideas I am going to suggest will perhaps illuminate some aspects of the first question. To cut a long story short, I defend here three hypotheses. The first is that, besides other interests, scientists have *epistemic* ones that can be reconstructed as the pursuit of a kind of "truthlikeness" (for example, the notion of truthlikeness proposed by Kuipers; see note 4 below). My second hypothesis is that scientists can engage in the negotiation of a

In: R. Festa, A. Aliseda and J. Peijnenburg (eds.), *Confirmation, Empirical Progress, and Truth Approximation* (*Poznań Studies in the Philosophy of the Sciences and the Humanities,* vol. 83), pp. 361-369. Amsterdam/New York, NY: Rodopi, 2005.

"methodological social contract," i.e., a set of norms indicating the circumstances under which a theory must be taken as better than its rivals; these norms act as the "rules of the game" of the research process, and tell scientists who must be deemed the winner of the game; some norms of this type are needed, because each scientist needs to know what 'winning' amounts to, if they are to become interested in playing the game at all. The last hypothesis will be that the choice of these norms is made by scientists "under the veil of ignorance," i.e., without having enough information about how each possible norm will affect the success of the theories each researcher will be proposing in the future. The main conclusion is that, under these circumstances, researchers will tend to prefer methodological norms which promote the truthlikeness of the theories which must be accepted according to those norms. This conclusion could be tested by studying whether the actual methodological norms used by scientists through history have been consistent, so to speak, with the maximization of verisimilitude.

Traditionally, philosophical explanations of science were developed under the tacit assumption that scientists disinterestedly pursued epistemic values, such as truth, certainty, generality, and so on. Even though sociologists in the school of Merton had put forward the fact that scientists were mainly motivated by other kinds of interests, this sociological school proposed the hypothesis that science was governed by an unwritten rule which forced scientists to disregard their personal or social motivations. This hypothesis presupposed that scientists were able to judge in an objective way which theory was the best solution to a given problem, and agreed in unanimously declaring that theory "the most appropriate one," even while many of them might have proposed alternative solutions. This utopian vision of the mechanism of scientific consensus has been challenged since the seventies by several new schools in the sociology of science, particularly by the two called "Strong Program" and "Ethnomethodology." According to these new radical schools, the role of non-epistemic interests – either social or personal – in the process of scientific research was absolutely determining. Through an overwhelming amount of empirical work, both in the history of science and in "laboratory studies," these scholars claimed to have shown that scientists tended to take their decisions motivated almost exclusively by this other kind of interests, and, as a conclusion, they advanced the thesis that "established" scientific knowledge did not as a matter of fact mirror the hidden structure of the world, but only the "all-too-human" struggles between scientists. Scientific consensus would thus not be the outcome of a match between theoretical hypotheses and empirical evidence, but the result of quasi-economic negotiations for the control of power in society and within scientific disciplines.

Of the two kinds of motivations these sociologists have assumed to explain scientists' behavior (i.e., interests rooted in roles and social classes, and interests related to status within the scientific profession), I think the second one better reflects the actual interests of individual scientists. In the first place, "wide" social interests are less apt to explain researchers' ordinary choices when the problems they try to solve have weak or uncertain social implications. In the second place, the relevant social groups are usually fewer than the number of competing solutions, so that it is impossible to match them one by one in order to explain a researcher's choice of a solution as the consequence of his belonging to a certain social group.

"Recognition" or "authority" can thus be the main goal of individual scientists, in the sense that, when they face a choice among several options, if one of them clearly leads to a higher degree of recognition, they will always choose this one. But it seems difficult to accept that this can be the *only* motivation for devoting one's life to scientific research: after all, "recognition" could be gained through many other activities, from politics to sports or to the arts, all of them more rewarding than science in terms of fame and income, and perhaps less demanding in terms of intellectual effort. If somebody has chosen to spend his youth among incomprehensible formulae and boring third-rate laboratory work, we may assume that he will at least find some pleasure in the acquisition of knowledge. So I propose to make the benevolent assumption that, in those cases where no option clearly entails an advantage in terms of recognition, a researcher will tend to base his choices of theory (or hypothesis, or description of facts) on the epistemic worth of the available options. This presupposes that each scientist is capable of establishing an (at least partial) ordering of the available options according to some set of "cognitive values," "epistemic preferences" or "epistemic utilities," but I do not go so far as to suppose that different researchers make necessarily the same ordering. Stated somewhat differently, I assume that scientists have some *informed beliefs* about the correctness or incorrectness of the propositions they handle during their work. After all, if researchers are able to obtain knowledge about which actions will cause them to reach a higher level of recognition, as radical sociologists of science easily assume, it would seem absurd to deny that they may also gain information about which theories are probably more correct, which empirical data are more relevant and which strategies of reasoning are logically sound. My assumption, hence, is simply that in those cases when choosing the most valuable option from the epistemic point of view does not diminish the expected level of recognition, this option will be chosen.

The question, hence, is what are the epistemic values of scientists, those which make them prefer *ceteris paribus* some theories, hypotheses or models to others. The sociology of science literature is not very helpful here, since it has either usually ignored this question, or it has just tried to show that epistemic preferences did not play any relevant role at all. Perhaps this question might be answered with the help of an opinion poll among scientists, but it would be difficult to decide exactly what should be asked; a particularly problematic issue about this poll would be to design the questions in a neutral way with respect to rival methodological or epistemological theories. Still more problematic would be the fact that scientists' cognitive preferences will probably be *tacit*, and it may be difficult for them to articulate in a coherent and illuminating way those factors which lead them to value a particular theory, experiment or model. On the other hand, I think we should avoid the arrogant stance of those philosophers who think that real scientists are not a reliable direct source of criteria in epistemological matters. Perhaps the average scientist is not very good at deriving *philosophical* implications from his own work, and perhaps he is as ignorant of the formal and conceptual complexities of scholastic philosophy of science as we often are about *his own discipline's* complexities. But, since scientists are our paradigmatic experts in the *production* of knowledge, it can hardly be denied that their practices will embody, so to speak, the best available criteria for determining what should count as "knowledge."

One common theme in the so-called "deconstructionist" approaches to the sociology of science is that "knowledge is negotiated." I do not deny it is. As a social institution, science is a "persuasion game," and in order to get recognition you have to make your colleagues accept that the hypotheses advanced by you are better than the rest… even better than those hypotheses which they themselves proposed! Agreement, hence, will hardly be an immediate outcome of a direct confrontation with an intersubjective experience; only after many rounds of "I will only accept that if you also accept this" moves will an almost unanimous agreement be reached. But it is also difficult to believe that the *full* negotiation process of a scientific fact is reducible to that kind of exchange, for in science it is compulsory to offer a *justification* of whatever proposition a researcher accepts or presents as being plausible; many readers of the sociological literature about the "rhetoric" of science may legitimately ask why that "rhetoric" has any force at all, why does each scientist not simply ignore *all* his colleagues' arguments and stubbornly reject any hypotheses proposed by a "rival." I guess that a more verisimilar account of the process will show instead, that what each scientist tries to "negotiate" is the coherence of a proposition with the *criteria of acceptance* shared by his colleagues. This entails that if, during the negotiation of a fact, you have employed a type of argument, or other propositions as premises, so as to "force" your colleagues to accept a hypothesis, you will be constrained to accept in the future the validity of that type of argument, or the rightness of those propositions. Otherwise, those colleagues you happened to persuade by means of that strategy could reverse their former decisions and reject your hypotheses, if they realize that you do not honor the very arguments you had used to persuade them.

As long as the decision of accepting a given fact or model (probably under the pressure of negotiation with colleagues, both rivals and collaborators) is constrained by the necessity of supporting that decision with *reasons* which are coherent with the kinds of reasons one has employed in other arguments, a fundamental point in the analysis of "scientific negotiations" must be why certain *types of reasons* are accepted by scientists, especially the types of reasons which serve as justifications of the use of other (lower-level) reasons. If we call these "higher-level" reasons *methodological norms*, then our question is just why certain methodological norms are used within a scientific community, instead of other alternative sets of criteria. I plainly accept that even these norms can be a matter of "negotiation," but we must not forget that the requirement of coherence entails that, if a scientist has acquired through former "negotiations" the compromise of abiding by his decisions to accept certain criteria, it is possible that the future application of those criteria will force him, for example, to reject a hypothesis he himself had proposed. So, as long as *reasons* are used in negotiation processes – reasons whose domain of application is necessarily wider than the negotiation of a particular fact or theory – it will be *uncertain* for a scientist whether in other cases it will be still favorable for him (i.e., for the acceptability of his own hypotheses) that those reasons are widely accepted. Or, stated somewhat differently, when a researcher decides to accept or to contest a given methodological norm, it is very difficult for him to make an informed estimate of how much support his own theories will receive from that norm "in the long run."

If my argument of the last two paragraphs has some plausibility, it follows that the decision of accepting or rejecting a given methodological norm must be made *under a* "*veil of ignorance*," in the sense that the personal interests of scientists will hardly serve them as a guide in their choice. This does not entail that those interests actually play no role in the negotiation about "proper methods," but we can guess that their influence will tend to be rather blind, if it happens to exists at all; i.e., even if a group of scientists accepts a methodological norm *because they believe* that it will help them to fulfil their professional aspirations, it is equally likely that the actual effect of the norm will be to undermine those aspirations. Under these circumstances, the only reasonable option for scientists is to base their choice of methodological norms on their epistemic preferences referred to above, since these preferences will allow then to make a much easier, direct evaluation of those norms. So, the wider the applicability of a norm, the more exclusively based on epistemic reasons alone it is likely to be.

My suggestion is, then, to look for those methodological norms which are more prevalent in the history of science, and to use them as data to be explained by a hypothesis about the nature of the scientists' epistemic preferences. This strategy would allow us to take the theories about the aim of science not – or not only – as metaphysical exercises, but also as *empirical hypotheses*, which could be tested against the history of science. The question is, hence, *what the epistemic preferences of scientists can be if they have led to the choice of the methodological norms observed in the practice of science?* Nevertheless, before arguing in favor of a particular hypothesis about those epistemic preferences, it is important to clarify some delicate points:

a) In the first place, there is probably no such thing as "the" epistemic preferences of scientists, for different scientists can have different preferences, and even one and the same scientist can change his preferences from time to time, or from context to context. The very concept of a "negotiation" applied to the choice of a set of methodological norms entails that the rules chosen will not necessarily correspond to the optimum choice of every scientist; instead, it can resemble a process of bargaining in which everyone agrees to accept something less than their optimum, in exchange for concessions made by the other parties. In particular, the chosen methodological norms may be different in different scientific communities or disciplines, as well as they may vary in time. So, we might end up with the conclusion that the best explanation of actual methodological practices is a *combination* of different epistemic preferences. This does not preclude that simpler explanations will be prefered *ceteris paribus*.

b) In the second place, the hypotheses taken under consideration should not refer to epistemic utilities which are too complicated from the formal point of view. The strategy defended here is, to repeat, that actual scientific practices are our best criteria to look for what constitutes "objective knowledge," and that these practices tend to reflect the cognitive intuitions of the "experts" in scientific research. It is hardly believable that these intuitions need to be reconstructed by means of excessively intricate epistemic utility functions, particularly when the functions are so complicated that no relevant, easily applicable methodological norms can be derived from them.

c) In the third place, our strategy suggests we should inspect scientific practice in order to look for "negotiations" about methodological norms, rather than about facts or theories. Most empirical reports from historians and sociologists of science refer to the second kind of negotiation, where conflicting methodological criteria are *used*, rather than *discussed*; the conclusion of many case studies is that a researcher or group of researchers managed to "impose" a new method, but usually it is left unexplained why the *other* scientists accept that "imposition" at all, if it goes against their own interests. So, either negotiations on method are not frequent, or they have tended to be neglected in the study of historical cases, or probably both things are partially true. In fact, methodological norms can often be "negotiated" in an absolutely tacit way: as long as they constitute the "rules of the game" of a scientific discipline or subdiscipline, they determine the rewards a researcher could expect to have if he decided to become a member of it; so researchers can negotiate the rules by "voting with their feet," i.e., by going to those fields of research in which, among other things, the norms are most favorable from their own point of view.

d) In the fourth place, the empirical character of this strategy does not preclude a *normative* use of it, in a "Lakatosian" sense: once a certain hypothesis about the epistemic preferences of scientists has been sufficiently "corroborated" (which entails, among other things, that there are no better available explanations of the methodological practice of scientists), we could use that hypothesis to show, e.g., that in some historical episodes the "optimum" methodological norms have not been followed, due to the influence of some factors; these factors might include the inability to reach an agreement about the norms or the presence of non-epistemic interests which suggested to researchers that other methodological norms could have been more favorable to them. On the other hand, even if an agreement about norms exists, some choices of models or facts made by some researchers may be contrary to the recommendations of the norms, especially if strong enough enforcement mechanisms fail to become established by the scientific community.

e) In the fifth and last place, and again from an evaluative point of view, a hypothesis about the epistemic preferences of scientists must not be identified with the thesis that those are the preferences they *should* have. Some philosophers may find logical, epistemological or ontological arguments to criticize the cognitive goals revealed by scientists in their methodological choices. But it is difficult to understand, in principle, how the generality of scientists – i.e., of society's experts in the production of knowledge – could be "mistaken" about the validity of their own goals. In this sense I think we should not dismiss an argument such as Alvin Goldman's against some critics of scientific realism, when he asserts that we should put into brackets any philosophical theory which, on the basis of complicated lucubrations on the concept of meaning, "proved" that the statement "there are unknown facts" has no logical or practical sense (Goldman 1986). On the other hand, knowledge advances thanks to the invention of new ideas that may look strange at first sight, and this is also true in the case of philosophy; so new epistemological points of view should nevertheless be welcomed and discussed, even if their discussion does not end up in their acceptance.

Going back to the question of what the epistemic preferences of scientists can be, I have argued elsewhere that some notions of verisimilitude can help to explain some of the most general methodological practices observed in science.[1] In particular, the approach developed by Theo Kuipers is very successful at giving a justification of the hypothetico-deductive method as an efficient mechanism for selecting theories closer and closer to the truth. Hence, under the perspective offered in this short paper, we can understand Kuipers' contribution to the theory of verisimilitude as an *explanation* of the *fact* that that method is as widely employed as it is: if the epistemic preferences of scientists are such that they consider a theory to be better than another just if the former is closer to the truth than the latter,[2] then they will tend to prefer the hypothetico-deductive method to any other system of rules, *were they given the chance of collectively choosing a norm about the method for comparison of theories they were going to use within their discipline*.

One possible challenge for those epistemologists who defend other kinds of cognitive utilities would be to justify that these other preferences explain, just as well as the theory of verisimilitude, the methodological norms *actually* adopted by scientists. Sociologists of science should also try to offer alternative explanations of the extreme popularity of the hypothetico-deductive method. Even if these better explanations were actually provided (which I do not discard *a priori*), the methodological approach to the theory of verisimilitude would have had the beneficial effect of promoting the kind of research which had led to those empirical and conceptual results.

---

[1] See, for example Zamora Bonilla (2000, pp. 321-35)

[2] One of the simplest definitions of truthlikeness proposed by Kuipers is the following: if X is the set of physically possible systems, theory A is closer to the truth than theory B if and only if $\mathrm{Mod}(B)\cap X \subseteq \mathrm{Mod}(A)\cap X$ (i.e., all "successes" of B are "successes" of A) and $\mathrm{Mod}(A)\cap \mathrm{Comp}(X) \subseteq \mathrm{Mod}(A)\cap \mathrm{Comp}(X)$ (i.e., all "mistakes" of A are "mistakes" of B). Here, a "success" is taken as a physical system rightly described, and a "mistake" as a physical system wrongly described. Kuipers shows that the hypothetico-deductive method is "efficient" for truth approximation by proving that, if A is *actually* more truthlike than B, then for any possible set of empirical data, A will always have more empirical success than B, and hence, if scientists follow the hypothetico-deductive method (which commands then to prefer those theories with more confirmed predictions), then it cannot be the case that a theory that is less verisimilar is prefered to a more verisimilar theory. See, for example, Kuipers (1992, pp. 299-341).

ACKNOWLEDGMENTS

Research for this paper has been made possible under Spanish Goverment's research projects PB98-0495-C08-01 and BFF2002-03656.

U.N.E.D.
Depto. Logica y Filosofia de la Ciencia
Ciudad Universitaria
28040 Madrid
Spain

REFERENCES

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

Kuipers, T.A.F. (1992). Naive and Refined Truth Approximation. *Synthese* **93**, 299-341.

Popper, K.R. (1972). The Aim of Science. In: *Objective Knowledge*, pp. 191-205. Oxford: Clarendon Press.

Zamora Bonilla, J.P. (2000). Truthlikeness, Rationality and Scientific Method. *Synthese* **122**, 321-335.

Theo A. F. Kuipers

ON BRIDGING PHILOSOPHY AND SOCIOLOGY OF SCIENCE
REPLY TO JESÚS ZAMORA BONILLA

There is a difficult relationship between present-day sociologists of science and social epistemologists, on the one hand, and "neo-classical" philosophers of science, on the other. Both parties have difficulty in taking each other seriously. Hope should be derived from those scholars who seriously try to build bridges. Of course, bridge builders have to start from somewhere and the most promising constructors with a philosophy of science background are in my view Alvin Goldman (1999), Ilkka Niiniluoto (1999), and, last but not least, Jesús Zamora Bonilla (2000). In the latter's contribution to this volume Zamora Bonilla continues his very specific project of clearly specifying a kind of research agenda for studying bridge issues, in critical response to Ilkka Kieseppä's reservations about a methodological role of the theory of verisimilitude and David Resnik's arguments against the explanation of scientific method by appeal to scientific aims. Some of his main points are the following. (1) Gaining "recognition" is the dominant personal motivation of scientists, followed by trying to serve epistemic values. (2) Epistemic values can be served by methodological norms. (3) The norms have to be chosen under a "veil of ignorance" regarding the fate of the theories that will be proposed by certain scientists and hence the recognition they will get from them. (4) Hence, the most common norms in practice will best serve the relevant epistemic values. (5) Conversely, an adequate epistemic theory should enable us to justify these norms. (6) The HD method is very popular among scientists and is favorable for truth approximation, at least when both are explicated along the lines of ICR or along related lines, as presented by Zamora Bonilla. (7) The theory of truth approximation even justifies the popularity of the HD method.

Zamora Bonilla concludes with:

> One possible challenge for those epistemologists who defend other kinds of cognitive utilities would be to justify that these other preferences just as well explain as the theory of verisimilitude the methodological norms *actually* adopted by scientists. Sociologists of

In: R. Festa, A. Aliseda and J. Peijnenburg (eds.), *Confirmation, Empirical Progress, and Truth Approximation* (*Poznań Studies in the Philosophy of the Sciences and the Humanities,* vol. 83), pp. 370-372. Amsterdam/New York, NY: Rodopi, 2005.

> science should also try to offer alternative explanations of the extreme popularity of the hypothetico-deductive method (p. 366).

I would like to accept the first challenge and make the second somewhat more precise. Before doing so, though, I quote a statement from SiS (pp. 349-50).

> To be sure, scientists not only aim at cognitive goals like empirical success or even the truth of their theories, but they also have social aims like recognition and power, and hence means to reach such aims. And although these goals frequently strengthen each other, [the existence of] such convergences by no means implies that the conscious pursuit of these social goals is good for science.

By arguing that epistemic values are subordinate to recognition and methodological norms subordinate to epistemic values, the latter on the basis of a veil of ignorance regarding the ultimately resultant recognition, Zamora Bonilla greatly relativized the possible negative effects of the conscious pursuit of recognition for the pursuit of epistemic values such as empirical success and truth.

To What Extent Are Instrumentalist Epistemic Values Sufficient?

A dominant line of argumentation in ICR is that the instrumentalist methodology, that is, HD evaluation of theories, is functional for truth approximation. Hence, that methodology serves the sophisticated realist cognitive values, and hence, conversely, these values can explain and justify the popularity of this methodology, to whit comparative HD evaluation. So far I agree with Zamora Bonilla. However, I would also claim that this methodology serves instrumentalist epistemic values, notably empirical success, at least as well. At first sight, Zamora Bonilla seems to disagree, but this might be mere appearance. The reason is that his own explication of truth approximation (see Zamora Bonilla 2000, and references therein) is essentially of an epistemic nature. Like Niiniluoto's (1987) notion of "estimated truthlikeness," it is not an objective notion. However, unlike Niiniluoto's notion, that of Zamora Bonilla is not based on an objective one. On the other hand, like my objective explication, and in contrast to Niiniluoto's explication, Bonilla's explication straightforwardly supports HD evaluation. Hence, the question is whether Bonilla's explication goes further than instrumentalist epistemic purposes. If so, my claim would be that even his explication is more than strictly necessary for explaining and justifying HD evaluation. However, this is not the occasion to investigate this in detail.

For the moment the interesting question remains whether there are other reasons to favor the (constructive) realist epistemology relative to the instrumentalist one. In ICR I have given two such reasons, one of a long-term and one of a short-term nature. Only the realist can make sense of the long-term dynamics in science, practiced by instrumentalists and realists, in which theoretical terms become observation terms, viz., by accepting the relevant theories as the (strongest relevant) truth. This general outlook enables the realist to relativize for the short term a counterexample to a new theory that is an observational success of a competing theory by pointing out the possibility that the latter may be accidental (ICR, p. 237, p. 318) or, to use my favorite new term, that it may be "a lucky hit." In sum, although both epistemologies can explain and justify the popularity of the instrumentalist method, only the realist can make sense of the regular occurrence of long-term extensions of the observational language and the occasional short-term phenomenon of downplaying successes of old theories.

The Proper Challenge to Sociologists Regarding Non-Falsificationist Behavior

None of this alters the fact that the suggested explanations-cum-justifications of HD evaluation provide an invitation to sociologists of science to offer alternative explanations of the popularity of HD evaluation. To be more precise, sociologists of science have shown convincingly that scientists frequently demonstrate non-falsificationist behavior. However, they have been inclined to look for "social" explanations for that type of behavior, whereas in the view of Zamora Bonilla and myself, straightforward cognitive explanations can be given. Certainly the most important reason is the relativization of the cognitive role of falsification in the process of (even purely observational) truth approximation. This amounts to the difference between HD testing and HD evaluation of theories. Moreover, both methods leave room for many sensible ways in which a *prima facie* counterexample of a favorite theory can be questioned as such. For both claims, see ICR, Section 5.2.3, or SiS, Section 7.3.3. Finally, there is the possibility of the lucky hit nature of successes of a competing theory, referred to above. Hence, in all these cases there are cognitive reasons to expect that non-falsificationist behavior may serve epistemic purposes. To be sure, and this is a major point made by Zamora Bonilla, on average this may well be useful for gaining recognition. Hence, in regard to non-falsificationist behavior, the proper challenge to sociologists is to look for cases that cannot be explained in this convergent way.

REFERENCES

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

Niiniluoto, I. (1987). *Truthlikeness*. Dordrecht: Reidel.

Niiniluoto, I. (1999). *Critical Scientific Realism*. Oxford: Oxford University Press.

Zamora Bonilla, J.P. (2000). Truthlikeness, Rationality and Scientific Method. *Synthese* **122**, 321-335.