



Epistemic divergence and the publicity of scientific methods

Gualtiero Piccinini

Department of Philosophy, Washington University, Campus Box 1073, One Brookings Dr., St Louis, MO 63130-4899, USA

Received 20 May 2002; received in revised form 24 November 2002

Abstract

Epistemic divergence occurs when different investigators give different answers to the same question using evidence-collecting methods that are not public. Without following the principle that scientific methods must be public, scientific communities risk epistemic divergence. I explicate the notion of public method and argue that, to avoid the risk of epistemic divergence, scientific communities should (and do) apply only methods that are public.

© 2003 Elsevier Ltd. All rights reserved.

Keywords: Epistemic divergence; Public method; Intersubjective test; Reliabilism; Method of possible cases

The activities of the sciences that are taught are things that can be seen and there is none that is not visible in one form or another.

Hippocrates¹

1. Introduction

Scientific statements must be intersubjectively testable. If evidence for a statement cannot be obtained by different investigators, then neither the evidence nor the statement are scientific. Classical defenses of this principle have been given by [Herbert](#)

E-mail address: gpiccini@artsci.wustl.edu (G. Piccinini).

¹ In *The Science of Medicine*, anciently attributed to Hippocrates.

Feigl (1953, p. 11), Carl Hempel (1952, p. 22), Immanuel Kant (1965, p. 645), and Karl Popper (1959, p. 44). In a stimulating paper, Alvin Goldman (1997) reformulates this venerable principle of scientific methodology as follows:

Publicity principle: any method for collecting scientific evidence must be public.

He argues that no notion of public method makes it plausible that public methods have better epistemic credentials than non-public ones; therefore, according to Goldman, the publicity principle should be rejected (*ibid.*, Sect. 3). For him, scientific statements need not be intersubjectively testable.

Goldman's challenge to this methodological pillar of modern science should be answered. By looking at motivations for the publicity principle and undesirable consequences of its rejection, this paper suggests a plausible explication of the notion of public method and good reasons to keep the publicity principle. In the next section, I'll describe Goldman's challenge in more detail. I'll then apply Goldman's proposal to a controversy over evidence-collecting methods and argue that some consequences of Goldman's proposal are intolerable; hence, we should keep the publicity principle. In the following section, I'll offer a modified definition of public method and defend it over Goldman's. Finally, I'll discuss different types of violations of the publicity principle and the epistemic problems they generate. That discussion shows how the publicity principle can be effectively applied to methodological disputes, reinforcing the conclusion that we should endorse it.

2. Goldman on public methods

The principal topic of this inquiry is evidence-producing methods used by scientists within empirical scientific practices—what scientists usually describe in the *methods* section of their reports. In this context, here is Goldman's definition of *public methods*:

Method M is a public (intersubjective) evidence-producing method if and only if:
 (A) two or more investigators can severally apply M to the same questions, and
 (B) if different investigators were to apply M to the same questions, M would always (or usually) generate the same answers (induce the same beliefs) in those investigators. (Goldman, 1997, p. 534)

A method that's not public will be called *private*.

Goldman notices that in (A) there is an ambiguity in the scope of *can*. According to a *strong* reading of (A), two investigators *actually* exist who can apply M to the same questions. This is implausible, says Goldman, for what happens if there is only one investigator? If a method is public when only two investigators remain, it should remain public after one of them dies. So, Goldman turns to a *weak* reading of (A), according to which it is *possible* that two investigators exist who apply M to the same questions. But then, he replies:

[W]hy would the mere possibility that other creatures in an accessible world might apply the same method to the same question confer higher epistemic credentials on a belief? I detect no intuitive plausibility in this idea . . . Why must publicity be a necessary condition for evidence production? (Goldman, 1997, p. 537)

Goldman doubts that a different definition of public method will add plausibility to the publicity principle (*ibid.*, p. 537), and he urges his readers to reject the principle (*ibid.*, p. 525).

Goldman proposes replacing the publicity principle with two requirements inspired by his reliabilist epistemology. Goldman's epistemological program is designed to account for knowledge in terms of the reliability of methods and processes by which beliefs are acquired, where a method or process is reliable if and only if it 'leads to truth a sufficiently high percent of the time' (Goldman, 1992, p. 129; see also Goldman, 1986, p. 27). Unless otherwise noted, I'll use the term 'reliability' in Goldman's sense. Epistemological reliabilism is not incompatible with the publicity principle: a reliabilist could still hold that scientific methods must be public. Goldman, however, wants to dispense with the publicity principle in favour of two alternative requirements, which methods must satisfy in order to be scientifically legitimate. First, a method must be reliable. Second, 'a method should also satisfy a certain "negative" constraint, viz., that there not be (undefeated) evidence of its unreliability' (Goldman, 1997, p. 543).

A method that satisfies Goldman's two requirements may or may not satisfy the publicity principle, and *vice versa*. One exception to this is that methods satisfying Goldman's first requirement also satisfy clause (B): if a method M_1 is reliable, and if different investigators apply M_1 to the same question, then M_1 will generally yield the same answer, namely the true one. Hence, if a method is reliable, it satisfies clause (B). But in general Goldman's requirements and the publicity principle are orthogonal to one another. On one hand, a method M_2 may be public but yield false results (violating Goldman's first requirement), for example, because the experimental design doesn't rule out certain defeating conditions or alternative explanations. If M_2 's unreliability is discovered, M_2 will violate both of Goldman's requirements without ceasing to be a public method. Another method M_3 might be deemed to be unreliable even though it is both public and reliable (violating Goldman's second but not his first requirement), for example, because the evidence that M_3 is unreliable is flawed. On the other hand, a method M_4 may be reliable but not public, for example, because only one individual endowed with unique cognitive abilities can apply M_4 (violating clause (A)). The fact that the publicity principle and Goldman's requirements are logically independent establishes that Goldman's requirements are a genuine alternative to the publicity principle.

For Goldman, then, a scientific method M need not be public, as long as M is reliable and investigators applying M have no evidence that M is *unreliable*. This proposal is not without consequences. In this paper, I will not dwell on reliabilism in epistemology: I will be as neutral as possible about the prospect of Goldman's general program, which, as noted above, is compatible with maintaining the publicity principle. I will discuss only Goldman's attempt to dispense with the publicity prin-

ciple in favor of his reliabilist requirements for legitimate scientific methods. I'll argue that the publicity principle helps us evaluate the legitimacy of scientific methods in a way that Goldman's proposed replacement does not. Hence, whether we are reliabilist epistemologists or not, we should keep the publicity principle.

3. Epistemic divergence and the publicity principle

A community of investigators that gives up the publicity principle faces the risk of *epistemic divergence*. This occurs when different investigators answer the same question in different ways using private methods for collecting evidence. Disagreement leads to controversy, and in scientific controversies, researchers routinely criticize each others' methods.

In the absence of epistemic divergence, controversy is not a problem. To be sure, lack of epistemic divergence does not mean that investigators will easily reach agreement. Disagreements may be very hard to settle, for instance because different theoretical commitments lead to different interpretations of the evidence. Also, agreement among investigators may or may not be good: if a community agrees on a false proposition, that may be epistemically undesirable. Avoiding epistemic divergence is no guarantee for agreement, and agreement is no guarantee for truth. But when methods are public, and hence there is no epistemic divergence, it is at least possible that disagreement and mutual criticism are a healthy component of the scientific dialectic, leading to improvements in methods and theories.

In epistemic divergence, disagreement is a very different matter. When methods are private, the parties in the dispute share no means to prove that a method is flawed—they have no common epistemic ground on which to resolve their disagreement. As long as investigators are in epistemic divergence, their controversies can never be settled. Lack of epistemic divergence is thus a necessary condition for disagreement to be resolvable, or for agreement to be challenged with hopes that a new agreement will be reached. Because of this, I take epistemic divergence to be undesirable.

To illustrate, let me report an instructive episode. During a hike in the beautiful Italian mountains, I found myself in an isolated, tiny village. Only two old brothers, whose names are Mino and Nino, live there—in two separate houses. While looking after their common herd of cows, these curious shepherds investigate scientific questions, their favorite being: what is the centre of the universe?

To answer their questions, Mino and Nino have different and—to our eyes—peculiar methods. Mino contemplates the movements of the cows in the fields. Then he relaxes and goes in a state of trance, where he can read the position of the centre of the universe from the paths his beloved cows follow while looking for fresh grass. The result is that the centre of the universe is his house. Of course, the cows go in slightly different directions every day, but every day Mino's result is the same. His brother Nino has a different method. By listening to his cows' *moos* for many years, he has learned their language—so well that he can ask them questions. When Nino asks them where the centre of the universe is, their answer is *his* house.

Sometimes, Mino and Nino compare the results of their inquiries; when I reached their village, they were in the midst of a discussion. They offered me some *grappa* and told me about their research. After I understood their disagreement, I sipped and asked Mino: ‘Is your method reliable?’

‘Of course,’ he answered, explaining that he applied it many times and it always produced identical results. No matter where the cows go, their aggregate movements indicate that the centre of the universe is his house. Moreover, his elaborate epistemological theory entails that his method is reliable. But Nino interrupted him, claiming that his method was as reliable as Mino’s, if not more; any cow he questions invariably gives the same answer: the universe is placed exactly around *his* house. He also has an epistemological theory, which is different from his brother’s theory and—it goes without saying—entails that his method is reliable.

To get out of the impasse, I asked Nino: ‘Why don’t you teach Mino to talk to cows, so that he will convince himself of what you say?’ The answer came quick: Nino tried to teach cow language to Mino over and over again, but Mino never learned. He appears to lack some special communicative ability that Nino has. When Nino was done, Mino spoke. He had tried many times to teach Nino his method of relaxation, but Nino never mastered it. Presumably, Nino is too tense. For a moment, I considered asking if either of them had any evidence that either of their two methods was *unreliable*, but I realized that I could guess their answers. So, I finished my *grappa*, said ‘Grazie, e arrivederci,’ and, inebriated, left them to their discussions.

Mino and Nino’s case makes me very reluctant to give up the publicity principle. From their respective points of view, they’ve strived to satisfy all constraints Goldman has imposed thus far on legitimate methods: they believe their methods are reliable and have no evidence of their *unreliability*. Their beliefs about the centre of the universe are mistaken, but from their epistemic position they can’t see their error. Even if they argue forever, Mino and Nino have no common ground on which to adjudicate their dispute. Their methods are private, and they are in epistemic divergence.

No matter what Mino and Nino say, their methods and results are far-fetched enough that they strike *us* as unreliable. If they *are* unreliable, then we may appeal to Goldman’s reliability requirement and deem those methods illegitimate. That’s an easy call here. If we could always stipulate which methods are reliable and which aren’t from our metaphysical armchair, we would never risk epistemic divergence. But we can’t stipulate whether *our own* methods are reliable. Just like Mino and Nino, the most investigators can offer their colleagues is their *belief* that their methods are reliable; whether that belief is justified cannot be established by armchair philosophizing. Investigators generate their beliefs in their methods’ reliability by using their own methods, the very methods whose reliability might be questioned. If their methods happen to be unreliable, then their beliefs that their methods are reliable, generated by those same unreliable methods, are mistaken. Just like Mino and Nino, investigators *can’t prove* that their methods yield truth ‘a sufficiently high percent of the time,’ i.e. that their methods are reliable in Goldman’s sense.

Assuming that Mino and Nino’s methods are unreliable, we can use a distinction introduced by Goldman to clarify what Mino and Nino are lacking. In Goldman’s

terms, a justification of a method is *strong* if the method *is* reliable, *weak* if the method *is not* reliable while the investigator *does not believe* it to be unreliable and *has no reliable means* to tell that it is unreliable (Goldman, 1992, p. 129–131). The distinction between weak and strong justification is important within epistemological reliabilism. According to reliabilism, a necessary condition for a belief to be knowledge is that it be produced by a reliable method or process, but reliabilism does not require the investigator to know whether her method or process is reliable. Reliabilists also recognize that when an investigator uses an unreliable method or process that *seems* reliable to her, without having reliable means to discover its unreliability, there is an important sense in which the investigator's beliefs are *blameless* even though the generating method or process is unreliable. To account for that sense of blamelessness, the reliabilist calls those beliefs weakly justified. By definition, an investigator with weak justification has no means to discover that her justification is weak; from the investigator's perspective, having weak or strong justification is epistemically indistinguishable. If Mino and Nino's methods are unreliable, they have merely weak justifications, but they can't realize that their justifications are weak or do anything about that weakness until they go beyond Goldman's requirements. Even though in principle Goldman's requirements would deem Mino and Nino's methods illegitimate, in practice those requirements are of no help to Mino and Nino.

So far, I assumed that Mino and Nino's methods are unreliable. Although this seems a reasonable assumption to us, it may not be strongly justified. For our assumption is based on our beliefs, which are generated by our own methods, which may or may not be reliable. Perhaps, aside from questions about the centre of the universe, Mino and Nino's methods are reliable after all. I didn't ask them how their methods answer other questions—questions that don't involve the centre of the universe. We can imagine that Mino and Nino's methods yield true answers to every question that crosses their minds, the only exception being the one about the centre of the universe. If that were the case, their methods would be remarkably reliable, and if so, Mino and Nino's beliefs would be strongly justified. But their epistemic divergence about the centre of the universe would be there all the same. Even with reliable methods, then, the risk of epistemic divergence remains. Strong justification (in Goldman's sense) for one's beliefs and methods, besides being epistemically indistinguishable from weak justification, does not rule out epistemic divergence.²

Without the publicity principle, we—like the two poor shepherds—risk epistemic divergence. To the extent that we want to minimize that risk, we should ensure that our methods are public.

4. A plausible notion of public method

To be entitled to the publicity principle, we must say what method publicity amounts to. Recall that, for Goldman, *M* is a public method if and only if (A) two

² However, other senses of *reliability* do play a role. For example, methods that satisfy clause (B) in the definition of public method are often called *reliable* in the scientific literature. In scientific practices, this reliability criterion does help avoid epistemic divergence (Piccinini, forthcoming).

or more investigators can apply M to the same questions, and (B) if different investigators applied M to the same questions, M would generate the same beliefs in those investigators. Goldman believes (A) should be read as requiring the possibility that two investigators apply M , but he also believes that the resulting publicity principle has no plausibility. To help see the plausibility, we need to refine and further disambiguate Goldman's definition.

The first question is not about the scope of *can*, but the quantifier ranging over *investigators*. Goldman takes it to be existential. But suppose Mino and Nino had two sisters: Lina, a very calm person, and Rina, an excellent communicator. Suppose Lina learns Mino's relaxation method and Rina learns the language of cows. This would satisfy Goldman's definition, but Mino and Nino's methods would still fail to satisfy our pretheoretical notion of public method. This is because now, instead of two individuals applying two different methods, there are two micro-communities of investigators neither one of which can apply the other community's method. Increasing the number of members in each community of investigators is not going to eliminate the epistemic divergence between the communities unless members of each community learn to use the other community's methods. For a method to be public, *any* investigator—any member of the relevant scientific community—must be able to apply it. To guarantee this in our definition of public method, we should put a universal quantifier in front of *investigators*.

Let's see what happens now to the scope of *can*. The scope ambiguity is fourfold, for there is another quantifier ranging over *questions*. Presumably, this is universal too.³ So, in the disambiguated formulation, the possibility operator might go either in front of everything, or between the two universal quantifiers, in either order, or after the universal quantifiers. Here are the four possible disambiguations of (A):

- (i) $\diamond (i) (q) (A(iMq))$
- (ii) $(i) \diamond (q) (A(iMq))$
- (iii) $(q) \diamond (i) (A(iMq))$
- (iv) $(q) (i) \diamond (A(iMq))$

A is a three-place predicate meaning '*... applies ... to ...*', i ranges over *investigators*, and q ranges over *questions*. (M ranges over *methods*, and its quantifier—here omitted—goes in front of the whole definition.) It seems onerous to require—as in (i)—that there is a possible world where all investigators apply M to all questions. It would also be quite demanding to ask—as in (ii)—that, for any investigator, there is a possible world where she applies M to all questions, and even—as in (iii)—that, for any question, there is a possible world where all investigators apply M to that question. I think the most reasonable formulation is that, for any question and any investigator, there is a possible world where the investigator applies M to that question. For present purposes, (iv) will do.

³ Of course, a method yields answers only when applied to *appropriate* questions, where what counts as an appropriate question depends on the method. I will presuppose this qualification from now on.

Choosing the best disambiguation, however, is less important than specifying the accessibility relation between worlds. In establishing what methods are public we cannot consider all possible worlds. Obviously, it is logically possible that Mino speaks to cows, or Nino applies Mino's relaxation method. Almost any method *could* be public, but not all are. To determine what methods are public at world w , we must grant accessibility only to physically possible worlds where investigators apply the cognitive capacities they have at w , although they are allowed to receive extra training, time, equipment, and funding. Both Mino and Nino have genuinely tried to teach and learn each other's methods. Since they have failed, we may reasonably conclude that, in any world where their cognitive capacities are the same, Mino cannot speak to cows and Nino cannot relax like Mino. As we expected, their methods are private.

Before settling on a definition of public method, a quick comment on clause (B). Clause (B) goes as follows: if different investigators applied M to the same questions, M would generate the same beliefs in those investigators.⁴ I think (B) is fine as it is, as long as we take Goldman's *beliefs* to mean *beliefs about what the results are*. Otherwise, understanding *beliefs* more broadly, the clause will rarely if ever be satisfied: different investigators applying the same method to answer the same question usually develop very different beliefs. As every experimental scientist knows, even the single application of one method to one question by one investigator generates different beliefs in different investigators. That's why in scientific papers the *results* section is distinct from the *discussion* section. How scientific communities reach agreement in the long run is a difficult issue, one I won't address here. What matters to the publicity of scientific methods is only that, *ideally*, different investigators answering the same questions by the same methods get the same results.⁵

If this is right, Goldman's definition can be amended as follows:

Definition. M is a public method if and only if:

- (A) any investigator can apply M to any question (formulation (iv)), and
- (B) if different investigators applied M to the same question, M would generate the same results.⁶

By the present account, the publicity principle is an epistemic norm meant to exclude from science all methods that, reliable as they may appear, should not be trusted

⁴ Like Goldman does, I'll ignore possible defeating conditions, including possible mistakes on the part of investigators, which could spoil the output produced by applying M .

⁵ No doubt, some experimentalists or observers are more talented and creative than others, and some manage to obtain reliable data in delicate conditions in which others failed before. This by no means shows that talented investigators follow methods that aren't public. For their data to be fully accepted by the community, it still must be possible for other investigators to replicate their results or to apply the same methods in similar conditions.

⁶ The definition is currently atemporal in the sense that whether M is public at time t does not depend on whether any individual has the capacity to apply M at time t . If one doesn't like this, the definition can be made relative to time. For M to be public at time t , it can be required that M be applicable by individuals who are investigators at time t .

because either not all investigators can apply them to all pertinent questions or, if all investigators can apply them to the same questions, the methods yield different answers.

This points to a problem deeper than that of the quantity of investigators—the problem, as it were, is one of their *quality*. Goldman does not discuss this, but any definition of public scientific method presupposes a viable notion of investigator, namely an answer to who should be included in, and who should be excluded from, the qualified set of investigators. The present definition of method publicity, like Goldman’s definition, inherits the vagueness of the notion of *investigator*. This is not a defect of our definitions, but a call for a sharper understanding of how we use this important notion; I can only make a few suggestions here, and will by no means exhaust the subject.

To begin with, counting only accredited members of scientific communities as investigators would beg the question. One side of the problem concerns which cognitive capacities are legitimately used by an alleged investigator. Sadly, a blind person will probably not make a good astronomer, no matter how hard she tries. Nevertheless, it seems fair that eyesight is legitimately used in science. But then, an occultist might suggest that it’s just as fair, for those who claim magical intuition, to maintain that, in respect to that faculty, the rest of us are lacking. So, we need to draw the line between legitimate and illegitimate (alleged) cognitive capacities. The other side of the problem is that many people, cognitively gifted as they may be, lack the specific disciplinary training—not to mention time, equipment, and funding—necessary for applying the most sophisticated scientific methods. Scientists themselves know only some of their discipline’s methods; for the most part, they are in no position to try methods used by their colleagues down the hall, let alone researchers from other departments. To be a scientist, it is enough to learn and apply *some* methods—not all. So, in practice, most people cannot apply most scientific methods. This doesn’t seem to make scientific methods private, but we need to explain why.

Before proceeding, we should restrict our attention to *human* investigators, and consequently to methods that are public or private for human investigators. This is because non-human creatures might have methods that are public for them but non-public for us. If bats could do science, they should be allowed to employ echolocation. Other non-human investigators might need no methods at all: divine science, acquired by direct intuition of the truth, has no need for a publicity principle. Since not all creatures have the same publicity constraints, the notion of investigator relevant to human science is that of a human investigator. From now on, I will focus on human investigators.

In attacking the problem of the quality of investigators, I assume that human beings normally share some cognitive capacities and can communicate their output to one another. The sense of ‘normal’ being employed is not a moral but a biological one. ‘Normal capacities’ are those that a Martian biologist would say humans have, where ‘normal’ is shorthand for a set of biological conditions under which humans develop those capacities.⁷ Capacities normally shared by humans include perceptual

⁷ This way of putting it is due to Paul Griffiths.

capacities such as vision, audition, etc., reasoning capacities such as generating or following a syllogism, and motor skills that might be needed to manipulate experimental apparatus. Assuming that humans normally share at least some capacities and communicate with each other, it is possible for them to demonstrate that some of their capacities have the same outputs. For example, they can point their eyes in the same direction and ask each other whether they see the same things. By doing so, they can agree that they do share certain cognitive capacities and that those capacities yield similar outputs under similar conditions.

This won't always be straightforward: there are blind people, and colourblind people, and people with many kinds of cognitive deficits. But the cognitive capacities that humans normally share can be used to demonstrate that some individuals have cognitive deficits. Usually blind people can be convinced that they are blind, colourblind people can be convinced that they are colourblind, etc. Certainly people with good vision have no difficulty convincing themselves that someone has poor vision, or no vision at all. The same holds, *mutatis mutandis*, for other cognitive capacities normally shared by humans. These capacities are legitimately used in science because they are normally shared by humans, and the fact that they are normally shared can be demonstrated without begging the question of method publicity by comparing their output under similar circumstances.

The difference between legitimate and occult cognitive capacities is that lacking magical powers and the like cannot be traced to specific events or conditions that can be individuated and investigated in their own right by means of capacities normally shared by humans, like blindness can. If a person sees poorly, this can be demonstrated independently of the use of eyesight in scientific methods. Furthermore, sight comes in degrees that can be compared to one another, and sight deficits can be traced to physiological conditions. Magical intuition and other occult capacities have none of these features. So, scientific investigators should be allowed to employ cognitive capacities that are normally shared by humans, and no cognitive capacities that are not normally shared by humans.

The reach and power of human cognitive capacities is not fixed once and for all. Human capacities can be refined by experience, education, and training. This is why parents try to educate their children, why schools exist, and why becoming a scientist requires years of training. We can trace, describe, and evaluate the painstaking learning process scientists go through. We can point at their textbooks, test scores, and reports. We can attend their classes, visit their laboratories, and join their meetings. When someone is left behind in a class despite her best efforts, this is not because of an element of privacy in the training of scientists, but perhaps because she is a slow learner. And whether one is a fast or slow learner can be demonstrated independently of one's success at becoming a scientist. The existence of records of the learning process, and independent evidence that some learners are faster than others, is the mark that a method can be taught.

Perhaps, one day human cognitive capacities will be improved by applying artificial devices to our brains, or by evolving new cognitive capacities, or by genetic engineering. This would not introduce elements of privacy in human cognitive capacities. If humans acquire new cognitive capacities, those new capacities should

be demonstrated and studied by means of other humanly shared cognitive capacities, in the same way that humans study special capacities of animals, such as echolocation in bats.

In conclusion, anyone who possesses cognitive capacities normally shared by humans, and refines and then uses those capacities to gain knowledge, should count as an investigator. Given all the necessary time and resources, she should be capable of applying any method that purports to be public. If cognitively skilled people try to learn a method by taking all the necessary time and doing their best, but still don't make it, then that method isn't public.

None of this should suggest that determining exactly what human cognitive capacities are—and what they can be used for—is an easy task. It can be very difficult. In some cases, it may take years of research and disputes to sort out delusions from genuine observations. The history of science is also the history of humans refining their knowledge of what they can learn about nature by the judicious use of their cognitive capacities. The point is not that it's easy to know what we can observe, but that it's possible. And to the extent that it's possible, it can be done by studying human cognitive capacities using public methods of inquiry. To find out more about human cognitive capacities and their powers, and to explain them, there are sciences of mind and brain. Ideally, these disciplines should follow the same methodological standards of other sciences, including employing public methods for collecting evidence.

Mature scientific communities are not deterred by the lack of complete knowledge about human cognitive capacities and what they allow humans to observe. When there is a dispute about a method, they implicitly appeal to the publicity principle and compare notes until they agree on what the method can be said to accomplish independently of which investigator applies it. For instance, when some researchers announced that they produced 'cold fusion' following a certain method, no other scientist concluded that their method was private but reliable. When others couldn't replicate their results, eventually the whole community rejected the original putative findings. By studying this practice, we may improve our understanding of method publicity. At any rate, in evaluating scientific methods, we should try to prevent the risk of epistemic divergence.

It may be objected that the above restrictions on who counts as legitimate investigators defeat the purpose of the universal quantifier in front of investigators (in the definition of public method). Why say that all investigators must be able to apply a method if then we count only some people as investigators? Wouldn't it be more honest to say from the start that only some people must be able to apply a method for it to be public? As long as we agree that we can't expect every human being to apply every method, these questions boil down to a question of semantic preference. There are two possible answers. Either not all humans are investigators and all investigators must be able to use a method for it to be public, or all humans are investigators but not all investigators must be able to use a method for it to be public. Either wording makes the point. I adopted the former because it reflects the common usage of 'scientist' as an honorific title that fully applies only to people with special credentials. For any method, only some people have all the (independently and pub-

licly demonstrable) capacities that, given the appropriate training, are necessary to apply that method. Those people form the relevant class of investigators, and all of them must be able to apply the method for it to be public.

5. Types of divergence

Given the present definition, there are three ways in which a method can be private. Correspondingly, there are three primitive types of epistemic divergence.

Type one: Every time different investigators apply the same method to the same question and obtain *incompatible results*, clause (B) is violated. A remediable version of this situation occurs when an investigator makes a mistake in applying a method, for instance because she doesn't control for some variable. In this case, other investigators applying the same method while controlling for that variable will get different results. However, courtesy of the publicity principle, scientists constantly look for flaws in the design or execution of experiments, and they settle this type of controversy by pointing out the error in the original experiment. Without the publicity principle, any scientist would be entitled to hold on to her own application of a method even when her results are different from those of her colleagues.

A good example of a non-scientific method that generates this first type of divergence is the method of possible cases—a favorite of some contemporary philosophers (Jackson, 1998). Investigators applying this method gather evidence by consulting their *a priori* intuitions about what is possible and impossible. Regrettably, different philosophers tend to have conflicting *a priori* intuitions. Suppose there are two philosophers, *A* and *B*. *A* conceives of a possible world *w* where *P* is true, and the truth of *P* at *w* entails thesis *T* at our world. *B* conceives of a possible world *w'* where *Q* is true, and the truth of *Q* at *w'* entails *not-T* at our world. Is *T* true or false at our world? Imagining more possible worlds isn't going to help, because the same problem arises for the newly imagined worlds. *A* and *B* disagree because they are trying to determine *T*'s truth value on the grounds of their intuitions, and their intuitions entail opposite conclusions. The two philosophers are in epistemic divergence of the first type, and the method of possible cases gives no remedy. Since the method of possible cases relies on *private* intuitions of what worlds are possible, it can generate epistemic divergence. To the extent that we want to avoid epistemic divergence, we should avoid the method of possible cases.

Type two: When a question *q* and a method *M* are such that *only some* investigators can apply *M* to answer *q*, clause (A) is violated. The result is that among the investigators who use *M*, only some will obtain an answer to *q* while others will be left with no answer (unless they use a method different from *M*). The clash between those with and those without answer is a peculiar form of epistemic divergence, but it's divergence nonetheless. An extreme case of this sort is introspection as construed by Goldman: 'Introspection is presumably a method that is applied "directly" to one's own mental condition and issues in beliefs about that condition' (Goldman, 1997, p. 532). For a question about someone's mental states, only the person whose mental states are in question can answer by introspecting. Moreover, Goldman

believes that contemporary psychologists rely on the method of introspection, in the sense that they take data collected by each introspecting subject and test hypotheses against those data (*ibid.*, p. 533). Throughout his paper, Goldman gives the impression that his main argument against the publicity principle is based on the use of introspective reports in psychology, which, he maintains, violates the publicity principle. Given this putative conflict between current scientific practice and the publicity principle, Goldman argues that *if* the use of introspection in psychology is legitimate, *then* the publicity principle should be rejected (*ibid.*, pp. 525–526). But Goldman's appeal to introspection is circular: his evidence that the use of introspection in psychology *is* legitimate (a premise that's necessary to affirm the consequent of his *modus ponens*) is the rejection of the publicity principle (*ibid.*, p. 543). The circle is not vicious only because of Goldman's intuitive considerations against the publicity principle, which I reported in Section 1 above. Since for Goldman introspection's legitimacy is based on the rejection of the publicity principle, that legitimacy yields no additional weight to the rejection itself.

This raises the question of whether the use of introspection in psychology is legitimate. If Goldman is correct that psychologists rely on introspection in a way that violates the publicity principle, we should condemn its use as unscientific. Indeed, that's a popular philosophical response (Lyons, 1986, p. 150; Dennett, 1991, p. 70), but a premature one. Here, we have no room to focus on the interesting issues raised by introspection. Elsewhere, I argue that Goldman's argument about introspection is unsound in a philosophically interesting way. Contrary to what philosophers like Lyons and Dennett maintain, Goldman is correct that using introspection in psychology is legitimate. But contrary to what both Goldman and his opponents maintain, psychologists—when they generate scientific data from introspective reports—do follow methods that satisfy the publicity principle. That is, psychologists do not ask introspecting subjects to collect *data* by introspecting; rather, they record the subjects' reports, and then the psychologists themselves extract data from those reports by following public procedures that are analogous to those followed by other scientists in generating their results.⁸

Type three: Finally, there are cases like Mino and Nino's. In those cases, there are at least two groups of investigators. The two groups apply methods M_1 and M_2 , respectively, while neither group can apply the other group's method. Again, clause (A) is violated. Suppose that M_1 and M_2 generate mutually inconsistent answers to a question. Without the publicity principle, the two groups are condemned to epistemic divergence. It's easy to think of historical controversies that approximate this third type of divergence; philosophers, theologians, and mystics often argue against each other without accepting the validity of each other's methods (e.g. empiricists vs. rationalists, or logical positivists vs. existentialists). Until Goldman challenged the publicity principle, we had hoped the sciences were different from philosophy, theology, and other non-scientific disciplines. Not because scientific methods are more

⁸ For a more detailed discussion of introspection in psychology and method publicity, see Piccinini (forthcoming).

reliable, which—in Goldman’s sense of ‘reliability’—we have no means to prove, but because (among other reasons) scientific methods are public. One function of the publicity principle is to keep scientific communities from epistemic divergence.

The three primitive types of divergence can be combined to form complex ones. For example, assuming that gods are omniscient and don’t deceive, then, by definition, divine revelation is a reliable method for gathering evidence. Perhaps for this reason, many people consider divine revelation more trustworthy than any other method of inquiry. Yet, divine revelation has been absent from the sources of scientific evidence for quite some time. Why, if it seems reliable? One good reason is that unfortunately, purported divine revelations happen only to a few individuals in mysterious circumstances—others have no way to check that any revelations occurred, let alone that they were divine. Divine revelation is private at least in senses one and two and, if we count revelations from different gods as different methods, they’re private in sense three as well. It is hardly surprising, then, that alleged beneficiaries of divine revelations are sometimes found in fanatical epistemic divergence with each other. More often than not, self-proclaimed receivers of revelations give different answers to the same questions. They or their followers also deny the authenticity of one another’s revelations, while lacking any means independent of their own revelations to disqualify them. Disputes among them are notoriously in vain. Perhaps this is why they are tragically resolved, when they are, by slaughter more than by argument.

One more case deserves mentioning: Sometimes, conclusions obtained by private methods are opposite to conclusions obtained by public ones. One instance is the current debate between evolutionists and those creationists whose evidential base is a literal reading of the Bible. In this controversy, putative evidence yielded by divine revelation is confronted with the observations of contemporary biologists and geologists. This case might be called epistemic *semi*-divergence. Like full divergence, different methods lead different investigators to different conclusions about the same question—in this example, the origin of species. Another similarity is that opposing parties deny the validity of each other’s methods. However, the methods of evolutionary biology are fully public—creationists could learn and apply them, if they were so inclined. The problem is, creationists refuse to do so, while offering other investigators no methods to generate evidence for their creationist story.⁹ This is why the current debate over creationism—as opposed to the debate between Darwin and the creationist biologists of the late nineteenth century—is not a *scientific* controversy. Holding on to the publicity principle, we can denounce the privacy of divine revelation and reject the notion that creationism is science.¹⁰

⁹ An anonymous referee has suggested that the creationist method is public too, because it consists of reading the Bible literally, which every literate person can do. Of course, the Bible is as public as any scientific report. But the methods under discussion are not those for *reading* reports; they are those for acquiring the reported evidence in the first place. As the story goes, it took unique individuals and unrepeatable circumstances to collect the evidence reported in the Bible and other reports of divine revelations.

¹⁰ This conclusion does not apply to Intelligent Design Theory, where creationists attempt to undermine evolutionary theory by public methods.

A final question concerns whether it is ever possible to accept a private method because it is reliable, or reject it because unreliable.¹¹ For example, when someone claims to be reading the future by private means, we should be able to point at the quality of her predictive record to determine the reliability of the method, and decide its legitimacy in that way. The answer to this question is that determining a private method's reliability is possible only by using other methods—in our example, methods for determining whether the predictions to be tested are true or false. Depending on whether those other methods are public or private, two cases obtain. If those methods are private, then the parties are in epistemic divergence and no determination is possible. If those methods are public, that may lead to the acceptance of the private method based on its reliability. But its epistemic legitimacy will still derive from that of the public methods used to establish its reliability.¹² The newly-legitimate private method may then be used to establish the reliability of other private methods. In the absence of examples of this kind from the history of science—of which I know none—this possibility remains academic. Nevertheless, the possibility suggests that the publicity principle may be legitimately weakened into the following recursively formulated principle: scientific methods are legitimate if and only if they are either public or shown to be reliable by legitimate methods.

6. Conclusion

If scientific communities were to follow Goldman's advice of accepting private methods of inquiry, they would risk epistemic divergence. Different scientists would answer the same question differently by applying methods that other scientists cannot apply. They would also reject each others' private methods as unsound without being able to show their opponents what's wrong with them. Scientific controversies would never be resolved. If we dislike this prospect, we should endorse the principle that scientific methods must, indeed, be public. A method is public if and only if two conditions are satisfied: any investigator can apply the method to the same questions, and the method generates the same results regardless of who's applying it.

It took a long time of serious metaphysical and methodological disputations before the publicity principle on scientific methods was established and accepted by scientific communities—at least as a regulative ideal. We have good reasons to keep it.

Acknowledgements

Part of this paper was presented at Florida International University. Thanks to the members of the audience for their feedback. Carl Craver, Paul Griffiths, Brian Hep-

¹¹ I owe this question to an anonymous referee.

¹² Of course, at most we can establish reliability of a method—private or public—relative to the reliability of the other methods we use and whatever reliability they have. In other words, we still can't know whether our belief that a method is reliable is strongly or weakly justified.

burn, Peter Machamer, John Roberts, Andrea Scarantino, Becca Skloot, and several anonymous referees gave me extensive comments, from which this paper has greatly benefited.

References

- Dennett, D. (1991). *Consciousness explained*. Boston: Brown.
- Feigl, H. (1953). The scientific outlook: Naturalism and humanism. In H. Feigl, & M. Brodbeck (Eds.), *Readings in the philosophy of science* (pp. 8–18). New York: Appleton-Century-Crofts.
- Goldman, A. I. (1986). *Epistemology and cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A. I. (1992). *Liaisons: Philosophy meets the cognitive and social sciences*. Cambridge, MA: MIT Press.
- Goldman, A. I. (1997). Science, publicity, and consciousness. *Philosophy of Science*, 64, 525–545.
- Hempel, C. (1952). *Fundamentals of concept formation in empirical sciences*. Chicago: University of Chicago Press.
- Jackson, F. (1998). *From metaphysics to ethics: A defence of conceptual analysis*. Oxford: Clarendon Press.
- Kant, I. (1965). *Critique of pure reason* (N. Kemp Smith, Trans.). New York: St. Martin. First published 1781.
- Lyons, W. (1986). *The disappearance of introspection*. Cambridge, MA: MIT Press.
- Piccinini, G. (forthcoming). Data from introspective reports: upgrading from commonsense to science. *Journal of Consciousness Studies*.
- Popper, K. (1959). *The logic of scientific discovery*. New York: Basic Books.