Goldman on Perceptual Discrimination

Carlotta Pavese

3.24.14

Outline

Introduction to Goldman on perceptual discrimination

Goldman's Two Analyses

Why Goldman's Account Is So Novel

Outline

Introduction to Goldman on perceptual discrimination

Goldman's Two Analyses

Why Goldman's Account Is So Novel

Some examples

Recall the sheep-in-the-meadow case we were discussing before:

sheep-in-the-meadow

You're in the meadow, and you see a rock which looks to you like a sheep. So you say to yourself "There's a sheep in the meadow." In fact there is a sheep in the meadow (behind the rock, where you can't see it).

Some examples

▶ This seems to be a case where you have a justified true belief that there's a sheep in the meadow, that fails to be knowledge.

Some examples

- This seems to be a case where you have a justified true belief that there's a sheep in the meadow, that fails to be knowledge.
- Now, one salient feature of this case is that you can't really see the sheep. You just think that you do.

Some examples

- This seems to be a case where you have a justified true belief that there's a sheep in the meadow, that fails to be knowledge.
- Now, one salient feature of this case is that you can't really see the sheep. You just think that you do.
- The fact that there really is a sheep in the meadow, which you don't see, seems to be just a gratuitous accident.

Some examples

- This seems to be a case where you have a justified true belief that there's a sheep in the meadow, that fails to be knowledge.
- Now, one salient feature of this case is that you can't really see the sheep. You just think that you do.
- The fact that there really is a sheep in the meadow, which you don't see, seems to be just a gratuitous accident.
- It doesn't have anything to do with your belief or your evidence for your belief.

Another example

Another example to the same effect:

Evil Neuroscientist

You look at a shelf. An evil neuroscientist has some electrodes wired up to your brain, and he causes you to have visual experiences as if there were a clock on the shelf. As it turns out, there really is a clock on the shelf. (But you would still be having clock-experiences even if there weren't.)

Evil Neuroscientist

▶ In that case, it seems like you have a justified belief that there is a clock on the shelf, and it's true that there is a clock on the shelf, but you don't know that there is a clock there.

Evil Neuroscientist

- ▶ In that case, it seems like you have a justified belief that there is a clock on the shelf, and it's true that there is a clock on the shelf, but you don't know that there is a clock there.

- ▶ Relative to what you really see, it's just an accident that there happens to be a clock there..

Diagnosis

One solution to the Gettier Problem suggested by these cases is to say that you know that P iff

1. you truly believe that P and

The difficulty for this view is to explain what we mean by "it's being an accident that you are right about P."

Diagnosis

One solution to the Gettier Problem suggested by these cases
is to say that you know that P iff

1. you truly believe that P and
2. it's not an accident that you are right about P.

The difficulty for this view is to explain what we mean by "it's
being an accident that you are right about P."

Diagnosis

Cheater

Tommy comes across some evidence that his wife is sneaking around. She isn't at work when she's supposed to be, he finds some matches from a fancy nightclub in her car, and so on. When he asks his wife where she's been, she is evasive. This gives him some evidence for believing that his wife is having an affair. As a matter of fact, there is a simple explanation for all the evidence that Tommy has encountered: his wife is planning a surprise birthday party for him. But as it also turns out, his wife is having an affair, with her old boyfriend in Chicago. But she's very discreet about it, and so hasn't left any clues lying around.

Diagnosis

1. This is a case where Tommy has a true belief that his wife is unfaithful,

Diagnosis

1. This is a case where Tommy has a true belief that his wife is unfaithful,
2. and it's just an accident that his belief is true.

Diagnosis

1. This is a case where Tommy has a true belief that his wife is unfaithful,
2. and it's just an accident that his belief is true.
3. So far, so good. According to the present proposal, Tommy doesn't know that his wife is unfaithful.

Another case

Now, consider the following variation on the case:

A Variation on Cheater

Tommy goes on a business trip to Chicago, checks into his hotel, and goes up to his room. By accident, he gets off the elevator on the wrong floor, and opens the wrong door. But the door does open, and there is his wife, in bed with her old boyfriend.

Another case

In this case, too, Tommy has a true belief that his wife is
being unfaithful to him. And in this case, too, it seems to be
just an accident that Tommy has a true belief about this. But
in this case, we do want to say that Tommy knows that his
wife is being unfaithful.

Diagnosis

So the problem is to explain the difference between the kinds of "accidents" illustrated in the first case, which do block your true beliefs from counting as knowledge, and the kinds of "accidents" illustrated in the second case, which don't.

Another solution to Gettier cases?

Let's consider a different way of explaining these Gettier Cases. This new solution says that you know that P iff

1. you truly believe that P and you have evidence that P, and

Another solution to Gettier cases?

Let's consider a different way of explaining these Gettier Cases. This new solution says that you know that P iff

1. you truly believe that P and you have evidence that P, and
2. the fact that P is causally connected in the right way with your belief or your evidence. (Sometimes this solution omits the reference to "evidence" altogether, and just talks about causal connections between the fact that P and your belief.)

Does it work?

1. This sounds promising.

Does it work?

1. This sounds promising.
2. In the sheep-in-the-meadow case, the problem was that the real sheep played no role in causing your belief that there was a sheep in the meadow.

Does it work?

1. This sounds promising.

2. In the sheep-in-the-meadow case, the problem was that the real sheep played no role in causing your belief that there was a sheep in the meadow.

3. That's why we want to say it's just an accident that you get things right. In the Nogot/Haveit case, Haveit's Ford-ownership likewise plays no role in causing your belief.

Does it work?

1. When we consider questions "Why did so-and-so act that way?" there are two kinds of answers we can give.

Does it work?

1. When we consider questions "Why did so-and-so act that way?" there are two kinds of answers we can give.
2. One kind of answer is a causal answer.

Does it work?

1. When we consider questions "Why did so-and-so act that way?" there are two kinds of answers we can give.
2. One kind of answer is a causal answer.
   2.1 This is a case where there is some reason or causal explanation why the belief is true, or why the agent acted as he did, but these need not be reasons the agent is aware of.

Does it work?

1. When we consider questions "Why did so-and-so act that way?" there are two kinds of answers we can give.

2. One kind of answer is a causal answer.

   2.1 This is a case where there is some reason or causal explanation why the belief is true, or why the agent acted as he did, but these need not be reasons the agent is aware of.

3. A second kind of answer is a rationalizing answer.

Does it work?

1. When we consider questions "Why did so-and-so act that way?" there are two kinds of answers we can give.

2. One kind of answer is a causal answer.

   2.1 This is a case where there is some reason or causal explanation why the belief is true, or why the agent acted as he did, but these need not be reasons the agent is aware of.

3. A second kind of answer is a rationalizing answer.

   3.1 This has to be some reason that the agent has for believing what he does, or acting in the way he does.

First example of a rationalizing answer

Rationalizing, I

Suppose I say "Sue went to the fridge because the levels of
such-and-such chemicals in her bloodstream went down." Here
I could be telling either kind of story. If I just mean: the levels
of the chemicals in Sue's bloodstream went down, and that
made Sue feel a bit hungry, etc., then I'm just giving a causal
answer to the question "Why did Sue go to the fridge?" Sue
need not have any idea about the chemicals in her
bloodstream.

First example of a rationalizing answer

Rationalizing, I

Alternatively, suppose that Sue is very worried about the chemicals in her bloodstream, so she has a monitor plugged into her arm. When the levels of certain chemicals drop, Sue sees this right away on the monitor, and she goes to the fridge to get some food so that she can make the chemicals return to the normal level. There the story about the chemicals does report Sue's own reason for going to the fridge. So that is a rationalizing answer to the question "Why did Sue go to the fridge?"

Second example of a rationalizing answer

Rationalizing, II

Suppose you and your friend are climbing a mountain, and at one point you're holding the rope while your friend climbs up below you. As you stand there, you remember something your friend did that really pissed you off. A smoldering desire to kill your friend starts to grow in your mind. Of course, you would never act on that desire. But it shocks you and horrifies you that you would even think of it, and this makes you so nervous that you unintentionally lose your grip on the rope and your friend plummets to his death.

Second example of a rationalizing answer

Rationalizing, II

We can say "Your desire to kill your friend was one of the reasons why you let the rope slip." But this is only true if we're citing the causal reasons why you let the rope slip. It's not like you intentionally let go of the rope, in order to kill your friend.

The case of beliefs

Rationalizing beliefs

Suppose Kurt is intensely paranoid. He always believes someone is trying to kill him. You ask me why he believes this, and I say "Because he has such-and-such a defect in his brain." It would be most natural to understand me here as giving a causal explanation of Kurt's belief. But depending on the circumstances, I might also be giving Kurt's own reason for believing what he does. Perhaps Kurt has discovered the defect in his brain, and in his paranoid way, he takes it to be evidence that someone is tampering with his body. He takes it to be evidence that someone is trying to kill him. In that case, it would play a role in rationalizing Kurt's belief.

Causes/reasons

1. As this last case brings out, one and the same thing might be both a cause of Kurt's belief and something that Kurt takes to be evidence for his belief. So there is no rule that if something is a causal reason why you believe what you do, it can't also play a role in rationalizing your belief.)

Causes/reasons

1. As this last case brings out, one and the same thing might be both a cause of Kurt's belief and something that Kurt takes to be evidence for his belief. So there is no rule that if something is a causal reason why you believe what you do, it can't also play a role in rationalizing your belief.)

2. The present solution to the Gettier Problem says that what we have to add to a justified belief, to get knowledge, is not more in the way of rationalizations, but rather more in the way of the right sort of causal connections between your belief and the facts in the world that make it true.

Alvin Goldman
Cast of Characters

- Ph.D., Princeton 1965
- Board of Governors Professor of Philosophy and Cognitive Science Rutgers, The State University of New Jersey
- Author of 12 books, including *A Theory of Human Action* (1977) and *Epistemology and Cognition* (1986);

Goldman on perceptual knowledge

1. According to the Relevant Alternatives Theories we've been considering so far, eliminating an alternative Q, or "ruling that alternative out," turned on what sort of evidence you have against Q.

Goldman on perceptual knowledge

1. According to the Relevant Alternatives Theories we've been considering so far, eliminating an alternative Q, or "ruling that alternative out," turned on what sort of evidence you have against Q.

2. Goldman rejects this assumption.

Goldman on perceptual knowledge

1. According to the Relevant Alternatives Theories we've been considering so far, eliminating an alternative Q, or "ruling that alternative out," turned on what sort of evidence you have against Q.

2. Goldman rejects this assumption.

3. He says you don't need to have evidence against Q to eliminate it; it suffices if you're able to "perceptually discriminate" the situation in which P obtains from the one in which Q obtains.

Goldman on perceptual knowledge

The Twin case

Suppose there are two twins, Judy and Trudy. Suppose moreover that you're usually able to tell them apart. You don't know how you do it, but when you're confronted with one of the twins, you're usually right in your beliefs about which twin it is. You may not be aware of any evidence you use to tell the two apart. Instead, there are subtle differences in their faces and the way they walk that your brain picks up and processes, without any conscious intervention or assistance from you. The details of these cognitive mechanisms are hidden from you.

Goldman on perceptual knowledge

The Twin case
For all you can tell, you just end up with beliefs about which twin is which, and these beliefs tend to be reliable. They're right much more often than not. In such a case, you're able to discriminate between Judy and Trudy although you're not in conscious possession of any evidence that you use to tell them apart.

Goldman on perceptual knowledge

The view Goldman wants to defend says the following:

1. Whether you have perceptual knowledge that P is a matter of whether you can reliably discriminate P from its alternatives.

Goldman on perceptual knowledge

The view Goldman wants to defend says the following:

1. Whether you have perceptual knowledge that P is a matter of whether you can reliably discriminate P from its alternatives.

2. You don't have to discriminate P from all its alternatives, but only from the relevant alternatives.

Goldman on perceptual knowledge

The view Goldman wants to defend says the following:

1. Whether you have perceptual knowledge that P is a matter of whether you can reliably discriminate P from its alternatives.

2. You don't have to discriminate P from all its alternatives, but only from the relevant alternatives.

3. Reliably discriminating P from certain alternatives doesn't require you to have evidence telling against those alternatives.

Goldman on perceptual knowledge

The view Goldman wants to defend says the following:

1. Whether you have perceptual knowledge that P is a matter of whether you can reliably discriminate P from its alternatives.

2. You don't have to discriminate P from all its alternatives, but only from the relevant alternatives.

3. Reliably discriminating P from certain alternatives doesn't require you to have evidence telling against those alternatives.

4. It only requires you to have cognitive mechanisms which tend to produce correct beliefs about whether P or the alternatives obtain.

Goldman on perceptual knowledge

The view Goldman wants to defend says the following:

1. As in the Judy/Trudy case, you don't have to know how you discriminate P from its alternatives; nor do you have to know that the method you use is reliable.

Goldman on perceptual knowledge

The view Goldman wants to defend says the following:

1. As in the Judy/Trudy case, you don't have to know how you discriminate P from its alternatives; nor do you have to know that the method you use is reliable.

2. There just has to be some method, and it just has to be reliable.

Goldman on perceptual knowledge

The view Goldman wants to defend says the following:

1. As in the Judy/Trudy case, you don't have to know how you discriminate P from its alternatives; nor do you have to know that the method you use is reliable.

2. There just has to be some method, and it just has to be reliable.

3. So the basic idea of Goldman's view will be this: if you truly believe that P on the basis of your perceptual experiences, then you perceptually know that P, unless there's some alternative to P which is both relevant and which you can't reliably discriminate from P.

Outline

Introduction to Goldman on perceptual discrimination

Goldman's Two Analyses

Why Goldman's Account Is So Novel

Goldman's first analysis

Goldman's first attempt to capture his basic idea goes as
follows. S perceptually knows that P iff:

- ▶ P is true.

Goldman's first analysis

Goldman's first attempt to capture his basic idea goes as follows. S perceptually knows that P iff:

- ▶ P is true.
- ▶ S has a perceptual belief that P.

Goldman's first analysis

Goldman's first attempt to capture his basic idea goes as follows. S perceptually knows that P iff:

- P is true.
- S has a perceptual belief that P.
- There is no alternative Q which is both relevant and which is such that, if Q obtained, S would still have a perceptual belief that P.

Goldman's first analysis

- ▶ For example, suppose S looks out the window of his car and sees a barn.

Goldman's first analysis

- ▶ For example, suppose S looks out the window of his car and sees a barn.
- ▶ S forms the true belief that there's a barn.

Goldman's first analysis

▶ For example, suppose S looks out the window of his car and sees a barn.

▶ S forms the true belief that there's a barn.

▶ One alternative would be the hypothesis that it's just an empty field.

Goldman's first analysis

- ▶ For example, suppose S looks out the window of his car and sees a barn.
- ▶ S forms the true belief that there's a barn.
- ▶ One alternative would be the hypothesis that it's just an empty field.
- ▶ However, if that alternative obtained, then S would have different experiences and so he would no longer believe that there's a barn.

Goldman's first analysis

- ▶ For example, suppose S looks out the window of his car and sees a barn.
- ▶ S forms the true belief that there's a barn.
- ▶ One alternative would be the hypothesis that it's just an empty field.
- ▶ However, if that alternative obtained, then S would have different experiences and so he would no longer believe that there's a barn.
- ▶ Hence, S is able to discriminate this alternative from what believes to be true.

Goldman's first analysis

- Another alternative might be that it's not a barn but just a barn facade.

Goldman's first analysis

- ▶ Another alternative might be that it's not a barn but just a barn facade.
- ▶ If this alternative obtained, then we can suppose that S would have the same experiences he's now having, and so he would still believe that there's a barn.

Goldman's first analysis

▶ Another alternative might be that it's not a barn but just a barn facade.

▶ If this alternative obtained, then we can suppose that S would have the same experiences he's now having, and so he would still believe that there's a barn.

▶ Hence, S is not able to discriminate this alternative from what he believes to be true.

Goldman's first analysis

- ▶ Another alternative might be that it's not a barn but just a barn facade.
- ▶ If this alternative obtained, then we can suppose that S would have the same experiences he's now having, and so he would still believe that there's a barn.
- ▶ Hence, S is not able to discriminate this alternative from what he believes to be true.
- ▶ If this alternative is a relevant one, then S won't count as knowing that there's a barn.

Goldman's first analysis

- ▶ Another alternative might be that it's not a barn but just a barn facade.
- ▶ If this alternative obtained, then we can suppose that S would have the same experiences he's now having, and so he would still believe that there's a barn.
- ▶ Hence, S is not able to discriminate this alternative from what he believes to be true.
- ▶ If this alternative is a relevant one, then S won't count as knowing that there's a barn.
- ▶ However, if the only relevant alternatives are like the empty field alternative, then since S is able to discriminate them from what he believes to be true, his belief will count as knowledge.

Problems for Goldman's first analysis

A dachshund

Suppose S sees a dachshund, and that S is able to reliably
discriminate dachshunds from larger animals like German
Shepherds and wolves. However, we suppose that S is not able
to tell German Shepherds and wolves apart very well. Now
suppose that there are lots of wolves running around, so the
possibility that what S sees is a wolf is a relevant alternative to
its being a dachshund. Intuitively, S still knows that what he
sees is a dachshund, since he's able to reliably discriminate
dachshunds from wolves. So far, so good. But now suppose
that S looks at the dachshund and forms the true perceptual
belief, "That's a dog."

Problems for Goldman's first analysis

A dachshund
Does this belief count as knowledge? Intuitively, it would. The fact that S mistakes some dogs (German Shepherds) for wolves shouldn't show that S can't know that this dog (which looks nothing like a German Shepherd) is a dog. So the intuitively correct thing to say is that S knows that the creature he sees is a dog, even though there are some other dogs which he can't tell apart from wolves.

The problem

Goldman's first analysis doesn't deliver this result though.
Goldman's first analysis tells us that S's belief "That's a dog"
doesn't count as knowledge, since there's a relevant
alternative–the possibility that it's a wolf–and if that
alternative obtained, we can suppose that S would confuse the
wolf for a German Shepherd and so still hold the belief "That's
a dog." So according to the first analysis, S's doesn't know
the dachshund he sees to be a dog.

Problems for Goldman's first analysis

What's gone wrong?

► If the wolf alternative obtained, S would still believe "That's a dog" but he'd believe it on the basis of different experiences.

Problems for Goldman's first analysis

What's gone wrong?

► If the wolf alternative obtained, S would still believe "That's a dog" but he'd believe it on the basis of different experiences.

► Perhaps that's the problem.

Problems for Goldman's first analysis

What's gone wrong?

- ▶ If the wolf alternative obtained, S would still believe "That's a dog" but he'd believe it on the basis of different experiences.
- ▶ Perhaps that's the problem.
- ▶ The fact that S confuses wolves with some dogs, ones that look very different, ought not to show that S can't know that the dachshund he sees is a dog.

Problems for Goldman's first analysis

What's gone wrong?

- ▶ As things stand, Goldman requires there to be no
  alternative Q such that, if it obtained, S would continue
  to form the belief that P for any reason.

Problems for Goldman's first analysis

What's gone wrong?

- ▶ As things stand, Goldman requires there to be no alternative Q such that, if it obtained, S would continue to form the belief that P for any reason.

- ▶ That's what gets him in trouble with the dachshund case. If S were seeing a wolf instead, S would still believe "That's a dog" but he'd believe it on the basis of very different experiences.

Problems for Goldman's first analysis

What's gone wrong?

- ▶ As things stand, Goldman requires there to be no alternative Q such that, if it obtained, S would continue to form the belief that P for any reason.

- ▶ That's what gets him in trouble with the dachshund case. If S were seeing a wolf instead, S would still believe "That's a dog" but he'd believe it on the basis of very different experiences.

- ▶ Perhaps Goldman should say instead that there is no alternative Q such that, if it obtained, S would continue to form the belief that P by the same method, or on the basis of the same sorts of experiences..

Goldman's second analysis

S perceptually knows that P iff:

1. P is true.

Goldman's second analysis

S perceptually knows that P iff:

1. P is true.

2. The state of affairs P causes S to have experiences E.

Goldman's second analysis

S perceptually knows that P iff:

1. P is true.
2. The state of affairs P causes S to have experiences E.
3. On the basis of those experiences E, S has a perceptual belief that P.

Goldman's second analysis

S perceptually knows that P iff:

1. P is true.
2. The state of affairs P causes S to have experiences E.
3. On the basis of those experiences E, S has a perceptual belief that P.
4. There is no alternative Q which is both relevant and which is such that, if Q obtained, S would still have a perceptual belief that P formed via the same method as actually produced his belief that P on the basis of experiences E.

Perceptually relevant alternatives

Call a possibility "perceptually equivalent" to P for S just in case the experiences it would produce in S are exactly the same as the experiences that P produces. (Or, if the experiences differ, they do so in respects that are ignored by the mechanisms that produced S's belief that P.) With this notion, we can restate the last condition in Goldman's analysis as follows:
(4) There is no alternative Q which is both relevant and which is "perceptually equivalent" to P for S.

Goldman's second analysis

This second analysis solves the dachshund problem.

1. If the wolf-alternative obtained, S would still form the belief "That's a dog," but he wouldn't form it by the same method.

Goldman's second analysis

This second analysis solves the dachshund problem.

1. If the wolf-alternative obtained, S would still form the belief "That's a dog," but he wouldn't form it by the same method.

2. He would form it on the basis of very different experiences.

Goldman's second analysis

This second analysis solves the dachshund problem.

1. If the wolf-alternative obtained, S would still form the belief "That's a dog," but he wouldn't form it by the same method.

2. He would form it on the basis of very different experiences.

3. So we get the intuitively correct result: the relevance of the wolf-alternative doesn't prevent S from knowing the dachshund he sees to be a dog.

Outline

Introduction to Goldman on perceptual discrimination

Goldman's Two Analyses

Why Goldman's Account Is So Novel

Internalism versus externalism

1. Call an account of some epistemological state internalist
   when it says that the presence or absence of the state
   depends on facts which are "internally available" to you,
   that is, knowable on the basis of introspection and
   reflection.

Internalism versus externalism

1. Call an account of some epistemological state internalist when it says that the presence or absence of the state depends on facts which are "internally available" to you, that is, knowable on the basis of introspection and reflection.

2. Call an account of some epistemological state externalist when it says that the presence or absence of the state depends on facts which aren't "internally available" to you. There can be internalist or externalist accounts of knowledge, of justification, and of various related notions.

Internalism versus externalism

1. Because of its focus on what you can reliably discriminate rather than on what sorts of evidence you have, Goldman's view is very different from traditional accounts of knowledge.

Internalism versus externalism

1. Because of its focus on what you can reliably discriminate rather than on what sorts of evidence you have, Goldman's view is very different from traditional accounts of knowledge.

2. On Goldman's view, whether your true beliefs count as knowledge doesn't depend on what evidence you're conscious of having, or on any other facts about your belief that are open to your introspective, self-reflective consciousness.

Internalism versus externalism

1. Because of its focus on what you can reliably discriminate rather than on what sorts of evidence you have, Goldman's view is very different from traditional accounts of knowledge.

2. On Goldman's view, whether your true beliefs count as knowledge doesn't depend on what evidence you're conscious of having, or on any other facts about your belief that are open to your introspective, self-reflective consciousness.

3. It depends on how your beliefs were caused, and whether they were caused in a way that's reliable.

Internalism versus externalism

1. Now, everybody says that whether your beliefs count as knowledge depends in part on whether they're true; and whether your beliefs are true is something that might not be detectable just on the basis of introspection and reflection.

Internalism versus externalism

1. Now, everybody says that whether your beliefs count as knowledge depends in part on whether they're true; and whether your beliefs are true is something that might not be detectable just on the basis of introspection and reflection.

2. So to that extent, everybody has an externalist theory of knowledge. But traditionally, philosophers have mostly thought that truth is the only externalist component of knowledge. They have assumed that all of the other features which go towards making true beliefs count as knowledge are "internally available."

Why Goldman's account is novel

1. Goldman's view rejects this traditional assumption.

Why Goldman's account is novel

1. Goldman's view rejects this traditional assumption.

2. On Goldman's view, whether your true beliefs count as knowledge depends on facts about the beliefs (how they were caused, how reliable the mechanisms which caused them are, etc.) which aren't "internally available" to you.

Why Goldman's account is novel

1. Goldman's view rejects this traditional assumption.

2. On Goldman's view, whether your true beliefs count as knowledge depends on facts about the beliefs (how they were caused, how reliable the mechanisms which caused them are, etc.) which aren't "internally available" to you.

3. When people talk about externalist theories of knowledge, they're usually thinking of views like this one.

Internal duplicates

1. Suppose there's someone out there who's a psychological duplicate of you.

Internal duplicates

1. Suppose there's someone out there who's a psychological duplicate of you.

2. He has all the same beliefs, thoughts, experiences, and memories as you have.

Internal duplicates

1. Suppose there's someone out there who's a psychological duplicate of you.

2. He has all the same beliefs, thoughts, experiences, and memories as you have.

3. Everything which you can tell about yourself on the basis of introspection and reflection, is also true of him (and like you, he can tell that it's true of him on the basis of introspection and reflection). Call such a person an internal epistemic duplicate of you.

The novelty of Goldman's view

1. On the traditional accounts of knowledge, if you had a belief which counted as knowledge, and your duplicate had the same belief, and his belief was also true, then his belief would have to count as knowledge too.

The novelty of Goldman's view

1. On the traditional accounts of knowledge, if you had a belief which counted as knowledge, and your duplicate had the same belief, and his belief was also true, then his belief would have to count as knowledge too.

2. On Goldman's account, however, your true belief might count as knowledge, whereas your duplicate's true belief fails to count as knowledge, because his belief was not produced in the same reliable way that yours was.

The novelty of Goldman's view

1. On the traditional accounts of knowledge, if you had a belief which counted as knowledge, and your duplicate had the same belief, and his belief was also true, then his belief would have to count as knowledge too.

2. On Goldman's account, however, your true belief might count as knowledge, whereas your duplicate's true belief fails to count as knowledge, because his belief was not produced in the same reliable way that yours was.

3. This is a difference between you and your duplicate, but it's not an "internally available" difference.

The novelty of Goldman's view

1. On the traditional accounts of knowledge, if you had a belief which counted as knowledge, and your duplicate had the same belief, and his belief was also true, then his belief would have to count as knowledge too.

2. On Goldman's account, however, your true belief might count as knowledge, whereas your duplicate's true belief fails to count as knowledge, because his belief was not produced in the same reliable way that yours was.

3. This is a difference between you and your duplicate, but it's not an "internally available" difference.

4. So on this view, there are other external features besides the truth-value of your belief which can make a difference to whether you know.