



Externalist (non-evidentialist) theories

Who cares about epistemic duties? (Not these guys.)



Externalist theories

- Causal theory
- Reliabilism
 - Counterfactual theories (e.g. truth tracking)
 - Process reliabilism
- Engineering (design) standpoint
 - Proper functionalism
 - Authoritarianism

Warrant or Justification?

- In Feldman chapter 5, externalists are presented as analysing “justification” in external terms.
- Feldman even presents Plantinga as giving an account of justification!


“Plantinga talks about “warrant” rather than “justification” in the passage quoted. We will take these terms to be equivalent, although Plantinga **may** have been interested in analyzing a concept somewhat different from justification.”

- N.B. Plantinga himself says (*Warrant: The current debate*, p. 46):

“No amount of dutifulness, epistemic or otherwise, is sufficient for warrant ... what we need to see clearly and first of all is the vast difference between justification and warrant. The lesson to be learned is that these two are not merely uneasy bedfellows; they are worlds apart.”

Overview: Causal theory

- S knows that p iff the fact p is causally connected in an appropriate way with S's believing p .



Overview: Reliabilism

- Reliabilism says that a belief is justified (or “warranted”?) when it is *reliable* in some sense.
 - **Counterfactual reliabilist theories** (e.g. Nozick, Dretske) understand reliability in terms of facts like, “if P weren’t true, then the subject would not believe that P” (owing to the laws of physics, and the construction of the cognitive process)
 - **Process reliabilism** (e.g. Goldman) says that a belief is justified by virtue of the reliability of the process that produced the belief.

The “engineering standpoint”

- Whatever caused the formation of living organisms must have had a strong “bias” towards making systems that accomplish certain *functions* (e.g. self-reproduction, getting energy from the environment, sensation, and motion).
- Thus organisms are (on all accounts) “designed” in some sense to accomplish certain functions.
- Belief formation is a biological process, and is therefore designed to work in a certain way, to accomplish specific functions.
- For a belief to have high epistemic status (warrant?) the cognitive process that produces it must be working properly, according to its design plan.

The causal theory

C. S knows p iff S's belief in p is caused by the fact p.

C*. S knows p if and only if the fact p is causally connected *in an appropriate way* with S's believing p.

- “Appropriate causal connections” include:
 - Perception
 - Memory
 - (Why not just *any* causal connection?)
 - proper reconstructions of causal chains. (Why?)

Reasonable inferences about the past

- E.g. Smith returns home and finds a lot of sawdust and wood chips where a tree used to be.
 - Does he know that the tree was cut down?
- Scientific theories about events in the distant past (e.g. evolution) are inferred from data observable today.

Knowledge of likely *effects*?

- According to the *Standard View*, not all our knowledge is of causes of our experience. We often obtain knowledge by reasoning about the likely effects of what we have observed.
- E.g. we see a fire in the fireplace, and conclude that there is smoke coming from the chimney.
- But we can't see the roof from our present position, so *the smoke didn't cause our belief that there's smoke*. That's why we have to “reconstruct the causal chain”.

Problems with the causal theory

1. Knowing generalisations
 - What does “all metal conducts electricity” cause?
 2. Overdetermination cases
 - You don’t have to reconstruct the causal chain properly.
You just need compelling *evidence* that the claim is true!
 3. Identical twins case
 - This case seems to show that one’s reconstruction of the causal chain must be *justified*.
- Verdict on the causal theory?

Reliabilism

- Reliabilism says that a belief is justified (or “warranted”?) when it is *reliable* in some sense.
 - **Counterfactual reliabilist theories** (e.g. Nozick, Dretske) say things like, “if P weren’t true, then the subject would not believe that P” (owing to the laws of physics, and the construction of the cognitive process)
 - **Process reliabilism** (e.g. Goldman) says that a belief is justified by virtue of the reliability of the process that produced the belief.

Nozick and 'truth tracking'

TT. S knows p iff

(i) p is true,

(ii) S believes p ,

(iii) S 's attitude toward p tracks the truth value of p :

— If p were **not** true, S would **not** have believed p ; and

— if matters had been different in a way that p remained true, S would still have believed p .

• (How does this handle the Gettier cases?)

- Nogot/Havit
- Wolf/sheep
- Fake barns

- Why is the second counterfactual needed?
 - (if matters had been different in a way that p remained true, S would still have believed p .)

Nozick: “If someone floating in a tank oblivious to everything around him is given (by direct electrical and chemical stimulation of the brain) the belief that he is floating in a tank with his brain being stimulated, then ... he does not know that it is true.”

Does this belief (that he’s in a tank) track the truth?

$\neg p \Box \rightarrow \neg B(p)$ holds

$p \Box \rightarrow B(p)$ fails

Lucky knowledge?

- Black is hard at work in her office. From time to time she looks up from her desk and computer to stretch her neck. On one such occasion she happens to glance out the window toward the street. Just at that moment she sees a mugging on the street. She has a clear view of the event. She is a witness.
- Black knows that a mugging has just occurred. Does this belief *track the truth*, however?
- N.B. “when p is true, S does believe p ” is usually taken to mean that S believes p in all the “close” possible worlds where p is true.

Toward process reliabilism?

TT*. S knows p iff

- i.* p is true,
- ii.* S believes p ,
- iii.* S used **method M** to form the belief in p , and
- iv.* when S uses **method M** to form beliefs about p , S 's beliefs about p track the truth of p .

- This modification leads to the *generality problem* of deciding which “method” or “process” was active in producing the particular belief that p ,
- Does this even fix the problem of lucky knowledge?

Variant of the fake barns case

- Smith sees what look like barns of various colours, then points to a (real) red one and forms two beliefs:

That's a barn, and

That's a red barn

Do these beliefs track the truth?

- The first belief doesn't track the truth, but the second one does! (The fake barns cannot be coloured red.)
 - So only the second belief is knowledge?

Scepticism and externalism

- According to reliabilism, and externalism generally, having knowledge doesn't require you to be able to rule out sceptical scenarios, using the resources you have cognitive access to.
- Having knowledge may require that you “couldn't be wrong”, but in the *objective* sense of possible, not the subjective sense.
- E.g. Nozick's externalism has interesting consequences:
 - I know there is one hand here
 - I don't know that I'm not a brain in a vat

Verdict on the tracking theory?

Process Reliabilism

- The justificational status of a belief is a function of the reliability of the process or processes that cause it, where (as a first approximation) reliability consists in the tendency of a process to produce beliefs that are true rather than false.
- N.B. If this “tendency” to produce true beliefs is understood as a probability, then it must be an *objective* probability, not a subjective one.

What is a 'reliable' process?

- “One that always (or almost always) works, across a wide range of circumstances.”
- Thus there are two dimensions to reliability:
 - i. A high *probability* or frequency of giving correct results (e.g. “works in 99% of cases”)
 - ii. Robustness: the probability in part (i) applies to a *wide range* of circumstances.

Range

- A reliable process is robust, i.e. it is almost always right in a wide (though limited) *range* of circumstances. E.g. a voltmeter:
 - has a maximum voltage before it gets fried.
 - works only in a certain temperature range,
 - gets affected by ionizing radiation.
 - withstands ordinary knocks, but won't work after getting walloped by a sledge hammer.

Is “lucky” knowledge reliable?

- Consider the earlier case, where Black (as a result of luckily glancing up at just the right time) knows that a mugging has occurred.
- Was this belief produced by a reliable cognitive process?

Degrees of reliability

- Reliability is obviously a matter of degree. No process is perfectly reliable, and some reliable processes are more reliable than others.
- Is this a problem for reliabilism?
- Reliabilists might say this is fine, since justification is a matter of degree as well.
 - But if knowledge is all-or-nothing, defining $K = RTB$ is going to lead to more Gettier problems.

Goldman's formulation (R*)

- 1. Basic beliefs** (non-inferential beliefs) are justified to the extent that they result from reliable cognitive processes (e.g. perception, memory).
- 2. Non-basic beliefs** are justified to the extent that they are produced by a reliable inferential process, using justified beliefs as premises.

Problem solved for foundationalism?

- A powerful objection to foundationalism was that the whole notion of a “justified basic belief” is highly suspect.
- Feldman suggested that a JBB is a “proper response” to a sensory experience, and that “experiences themselves can be evidence”.
- These claims (arguably) strain evidentialism beyond the breaking point!
- Reliabilism (and most other externalist theories) have no problem accounting for JBBs.

Externalist account of basic beliefs

- According to externalism, there is no need to find a *reason* why your basic beliefs (e.g. perceptual beliefs) are likely to be true.
- (Bonjour's TIF objection is completely avoided.)
- Like Emma the dog, we can simply follow our instincts.

Objections to reliabilism

1. Reliability is not *necessary* for justification

- The Brain in a Vat (or Cartesian evil genius) case. These people are epistemically blameless, being diligent, careful and thorough. Hence their beliefs are justified. But they're not at all reliable.

2. Reliability is not *sufficient* for justification

- Clairvoyance cases. These genuine clairvoyants have reliable beliefs, but the beliefs are unjustified because they have no reason to think they're true.

3. Generality Problem

- Talk of "*the* process" that produced belief in p is fatally ambiguous. Reliability is defined only for process *types*, and each token process is an instance of many different types.

1. Evil genius case

“Imagine a group of people who live in a world controlled by a Cartesian evil genius ...

...The people in this position are, we may suppose, careful and thorough investigators. They accumulate large quantities of sensory evidence, formulate hypotheses and theories, subject their beliefs to careful experimental and observational tests, and so on ...

... Are the beliefs about their apparent world that the people in such a Cartesian demon world arrive at in these ways justified? ... From an intuitive standpoint, it seems hard (doesn't it?) to deny that they are.”

(BonJour, *Epistemology*, p. 228)

2. Clairvoyance cases

... doesn't it seem as though Norman is being thoroughly irrational and so is not in fact justified in confidently accepting beliefs on this sort of basis?

(Think about this question on your own. One way to develop the issue further is to ask whether Norman would be justified in acting on one of these beliefs if an urgent occasion should arise: perhaps someone is trying to contact the president on an urgent matter and asks Norman if he knows where to find him.

(BonJour, p. 231)

Variants of the clairvoyance case

- In the case mentioned, Norman has independent evidence that his clairvoyant belief, that the President is in NYC, is false. (TV news, etc.)
- Bonjour also considers cases where Norman has scientific evidence that clairvoyance is impossible, all alleged clairvoyants have been exposed as fakes, etc.
- Such evidence makes no difference at all to Norman's reliability, yet it certainly makes his belief that the President is in NYC less justified!

3. Generality problem

“The **Generality Problem** is the problem of specifying exactly which process it is whose reliability determines how justified your belief is.”
(Prior)

E.g. I look out the window, and believe “it’s raining”.
Which process formed this belief? Is it:

- the process of forming beliefs on the basis of *perception*
- the process of forming beliefs on the basis of *vision*
- the process of forming beliefs about the weather on the basis of looking out a window
- the process of forming a belief that it’s raining on the basis of seeing droplets splashing on the pavement
- etc.

Types or tokens?

- The generality problem takes it for granted that reliability values are defined only for process *types* rather than *tokens*.
- (After all, each token process produces either a true belief or a false one. There is no success rate!)
- Is this correct?

- (Most people now think that there are single-case chances.
- Can a reliable process be defined as one with a high chance of success, in that exact set of circumstances?)

Generality problem

1. If we define the process broadly, e.g. “vision”, then the problem is that *some visual beliefs are a lot more justified than others*. Yet on this view they’d all be equally justified.
2. To avoid (1) we define the process as narrowly as possible. But then there may be only one belief that the process ever produces (no two perceptions are identical) and so it’s meaningless to ask how often it produces true beliefs.
 - Appeal to single-case *chances*, rather than frequencies?
 - But since a reliable cognitive mechanism is also *robust*, in the sense that it works over a wide range of circumstances, the process must be defined broadly?

Generality Problem

Without some way of answering this question in a specific and nonarbitrary way, the reliabilist has not succeeded in offering a definite position at all, but only a general schema that there is apparently no nonarbitrary way to make more definite. Certainly some ways of specifying the relevant process are more natural than others; but the epistemological relevance of such naturalness is questionable, and even these more natural specifications are numerous enough to result in significantly differing degrees of reliability.

Though reliabilists have struggled with this problem, no solution has yet been found that even a majority of reliabilists find acceptable. (BonJour, p. 215)

4. The Range Problem

(The reliability of a machine is always limited to a certain range of circumstances, or possible environments.)

“The **Range Problem** is the problem of specifying *where* a process has to be reliable--in what range of possible environments?--in order for beliefs produced by it to count as justified.”

- I.e. Knowledge seems to require some range of reliability, surrounding the actual circumstances. *What defines this range?*

Verdict on process reliabilism?

- Does process reliabilism avoid the Gettier counter-examples? In the sheep case, for example, is there any unreliable process involved?
- If not, then that's a *great loss*. Externalism was supposed to be the radical solution to the Gettier problem. Other externalist theories have avoided Gettier.

The “engineering standpoint”

- Belief formation is a biological process, and is therefore designed to work in a certain way, to accomplish specific functions.
- For a belief to have high epistemic status (warrant?) the cognitive process that produces it must be working properly, according to its design plan.
- First proposed by Alvin Plantinga in *Warrant and Proper Function*, 1993.
- One of Plantinga’s claims is that the “gap filler”, that turns true belief into knowledge, is *warrant* rather than justification.
- $K = WTB$.

Plantinga's "proper functionalism"

Plantinga says that a belief, B, is warranted if:

- (1) the cognitive faculties involved in the production of B are functioning properly...;
- (2) your cognitive environment is sufficiently similar to the one for which your cognitive faculties are designed;
- (3) ... the design plan governing the production of the belief in question involves, as purpose or function, the production of true beliefs...; and
- (4) the design plan is a good one: that is, there is a high statistical or objective probability that a belief produced in accordance with the relevant segment of the design plan in that sort of environment is true.

Reliability isn't sufficient for warrant

Plantinga's "Epistemically Serendipitous Brain Lesion"

“Suppose that Sam suffers from a serious abnormality – a brain lesion, let’s say. This lesion wreaks havoc with Sam’s noetic structure, causing him to believe a variety of propositions, most of which are wildly false. It also causes him to believe, however, that he is suffering from a brain lesion. Further, Sam has no evidence at all that he is abnormal in this way, thinking of his unusual beliefs as resulting from an engagingly original turn of mind.”

Accidentally reliable processes?

- Plantinga stipulates that the brain lesion is a *reliable* producer of the belief that one has a brain lesion, yet it seems that the belief isn't knowledge because the brain isn't *designed* to produce this belief.
- Does that sound like the right response?
- In that case, having a reliable true belief isn't *sufficient* for knowledge.
- How does this compare to Bonjour's clairvoyance case?

Clairvoyant *by design*?

- What if Norman isn't just a reliable clairvoyant, but is *designed* to be one? (The other perceptual mechanisms are designed, after all.)
- What if Norman's clairvoyant beliefs are "clear and distinct", assuring him of their reality? (This is of course a design feature as well.)
- We don't think ordinary perceptual beliefs need independent verification. (They're basic.) Why then would clairvoyant beliefs need this?

Engineering and “assurance”

- “Assurance” here means a cognitively accessible indicator of epistemic well being.

E.g. according to Locke, one reason to trust the senses is:

... the assurance we have from our senses themselves that they don't err in what they tell us about the existence of things outside us when we are affected by them ...

I am certain that I am a thinking thing; but do I not therefore likewise know what is required to render me certain of a truth ? In this first knowledge, doubtless, there is nothing that gives me assurance of its truth except the clear and distinct perception of what I affirm, which would not indeed be sufficient to give me the assurance that what I say is true, if it could ever happen that anything I thus clearly and distinctly perceived should prove false.

Descartes, *Meditation III*.

(Apparently this clarity and distinctness is both cognitively accessible, and an infallible TIF.)

Engineering and “assurance”

- From an engineering perspective, assurance is an essential component of a cognitive system.
 - How is the “user” of the cognitive system going to distinguish between reliable products of the system and mere idle imaginings?
 - How is the user to be made aware of the fact that some beliefs are more reliable than others?
 - The engineering solution is for each belief to be accompanied by a phenomenological feel, in proportion to its reliability. The user has a duty to pay attention to this degree of assurance.

Engineering and “assurance”

- But how does the user know *which* phenomenal feels are TIFs? How do they know which spontaneous beliefs to trust?
- Perhaps the user is simply designed to *automatically* trust beliefs that are accompanied by assurance?
 - Humans feel strongly compelled to cease activities that cause pain, or otherwise “feel wrong”, and to continue activities that cause feelings of pleasure and well being.

- Plantinga assumes that when a person trusts and follows their feelings of epistemic well being, their beliefs are *justified*.
- (Like the person monitoring a nuclear reactor, relying on the dials. Or a Brain in a Vat. What more can they be expected to do?)
- But if the cognitive processes are malfunctioning, producing spurious assurance, then the belief isn't *warranted*.

- Plantinga's view (that accepting beliefs that are accompanied by assurance is justified) is probably right in most cases.
- For animals, it is surely right in all cases.
- Humans however can be aware of their own belief formation, and think critically about it. Arguably, this ability entails added epistemic duties in some cases.



**WITH GREAT
POWER
COMES GREAT
RESPONSIBILITY**

- Humans can be aware of their own belief formation, and think critically about it.
- Arguably, this ability entails added epistemic duties in some cases.

E.g. Stanislav Petrov saves the world

- In 1983 Petrov was the duty officer for a Soviet nuclear early-warning system when the system reported that 3-5 ICBMs had been launched from the United States.
- According to protocol, Petrov should have reported this attack to superiors. Had he done so, the USSR would probably have retaliated with its own nuclear weapons. ☹️
- Instead, Petrov judged that it was (probably) a false alarm, and **did not report it**. Why?
 - For the USA to launch so few missiles would be suicide
 - The automatic detection system was new
 - Ground radar found no corroborative evidence

Stanislav Petrov case

- Note that, in this case, Petrov had multiple sources of knowledge:
 - The automatic missile detection system
 - His background knowledge of nuclear conflict
 - Ground radar signals
- Given that they conflicted, he was (arguably) right to be sceptical of the missile warning.

Multiple cognitive mechanisms

- Like Petrov, every person has multiple cognitive faculties (vision, hearing, touch, reasoning, memory, etc.) which overlap in the sense that they can agree with or contradict each other.
- E.g. you see (or seem to see) a pink elephant, floating in the air.
- E.g. Norman has a clairvoyant experience, accompanied by assurance, yet science disavows clairvoyance.
- (What would a cognitive engineer design the agent to do in such cases?)

- In some cases at least, reasoning overrules perception.
- In other cases, perception overrules previously strong convictions. (Your friend had a leg amputated. You saw the stump. But later you see him again with two legs, and he says it grew back. Will you *ever* believe this happened?)
- In such clashes, what determines the winner?
- Presumably there is some faculty (reason?) that is meant to decide such things.
- In cases of conflict, a properly-functioning agent does *not* simply trust the assurance of the senses?

Epistemic duties for externalism?

- If a proper-functionalist takes this route, then they're apparently making some concessions to internalism, or agreeing with some internalist motivations.
- They're saying that (in some cases at least) one has a duty to doubt one's spontaneous beliefs, and verify them using other factors that are internally available (e.g. other beliefs).
- (Of course the overruling is handled by another "failsafe" process that is designed to do that.)

Is any compromise possible between internalism and externalism?

“We may begin by asking whether it is really as clear as I have in effect been assuming (and as those on both sides of this issue typically assume as well) that the internalist and the externalist views of justification are incompatible in a way that means that one must be simply right and the other simply wrong.

Some philosophers have in fact suggested that perhaps there are instead two (or even more) different conceptions of knowledge or justification, one (or more) of them internalist and one (or more) of them externalist: conceptions that simply address different issues and serve different purposes, and that are thus not in any meaningful sense competitors between which a choice must be made.”

BonJour, 215-216

Problem for the engineering standpoint

- What if the designer is incompetent?
 - An incompetent designer might make a cognitive mechanism that is accidentally reliable!
 - (There's a difference between functioning *as designed* and functioning *well*.)
- In fact there are only two candidate designers: the evolutionary process and God
- Plantinga himself is a theist, and thinks that the designer is God. Would that help?

Questions for Plantinga

1. Why is proper functioning the essential feature of a warranted belief (and hence of knowledge)?
 - E.g. how is proper functioning related to truth?
2. Bonjour's question: "Why should some external facts and not others be relevant to justification?"
 - Isn't proper functionalism *ad hoc*?