

Announcements and Such

- One Song — *Dub Resistance*
 - “Ridin” from *World Receiver*
- **Second Essay Topics Posted Today** (due 4/5):
 - I also list the topics at the end of this lecture
 - We can discuss these next time, if people want
- **Please attend lecture Thursday** (we’ve a visitor)
- Today: Part II of *The Analysis of Knowledge*
 - *Naturalist K*: Causal, reliabilist, truth-tracking
 - Prelude to Internalism vs Externalism
- Vanessa’s OH are 2:15–4 today, since she’s covering James’s sections on Thursday afternoon.

The Analysis of Knowledge II Naturalistic Accounts of Knowledge

- So far, we’ve been supposing that *justification* is one of the components of knowledge.
- As we have seen, though, “JXTB” theories run into some serious challenges (Gettier-Zagzebski style).
- Perhaps we should take a different approach entirely, which *doesn’t* require justification at all.
- Audi describes such approaches as “naturalistic” — they think of knowledge as a matter of *registering truth*, much as *thermometers* register *temperature*
- On such approaches, we think of agents as “belief generating devices”, and we focus on their *physical* and *causal* properties. *Justification plays no role*.
- Audi discusses two such approaches: *the causal theory*, and *reliabilism*. I will also discuss *truth-tracking* accounts (a 3rd) in detail (Audi does not)

The Analysis of Knowledge II The Causal Theory of Knowledge

- On the causal theory, a true belief constitutes knowledge if it is caused in an appropriate way.
- The paradigm cases favorable to the causal view are simple (good) perceptual cases. The green field before me causes me to (truly) believe it is there.
- This is a *direct* causal relation between the object of my true belief and the true belief in question.
- The causal theory can also ground knowledge about the future, but that will involve a different sort of (*indirect*, common-cause) causal structure.
- There can be *common causes* of (both) my belief and the future event in question. In such cases, the causal account *can* accommodate knowledge.
- What about *a priori* knowledge? Is *that causal*?

The Analysis of Knowledge II Knowledge as Reliably Grounded True Belief I

- The causal theory can run into trouble because it doesn’t require *reliable* belief-generating processes
 - Example: Jim’s being angry causes (*via* observation) Tom to believe (and *know*) Jim is angry, and Tom’s testimony causes me to believe (truly) that Jim is angry. But, Tom is highly unreliable in his reports about Jim’s emotions.
- *Reliabilism*: knowledge is *reliably grounded* true belief. On this account, we are to focus on the *reliability* of the *belief-generating process*.
- The causal theory has no good way to block this as a case of knowledge for me. A reliabilist will deny that this (true belief) constitutes knowledge.
- Reliabilism can also explain why perception, reason, *etc.*, yield knowledge *when they’re reliable*

The Analysis of Knowledge II Knowledge as Reliably Grounded True Belief II

- So, reliabilism can handle *a priori* knowledge too (in terms of *reliably reason-generated* beliefs)
- Crucial Question about Reliabilism: *How* reliable must a belief-generating process be in order for a true belief that it generates to count as *knowledge*?
- I'll return to this question again later on.
- There is another, even more problematic question we must ask about reliabilism: *how do we individuate "the belief-generating process"*?
- Under some descriptions of a process, a belief may be reliably generated, under other descriptions of a process, a belief may be unreliably generated.
- This is sometimes called the "specification" or "generality" problem. It's quite a difficult one.

The Analysis of Knowledge II Knowledge as Reliably Grounded True Belief III

- The reliability of, *e.g.*, vision varies greatly with the *conditions of observation*. So, it's misleading to say (simply) that vision is a *reliable* BGP.
- Moreover, even in the very same context, we can describe the same BGP in different (true) ways, leading to different assessments of reliability.
- This is related to the *non-monotonicity of probability*.
- Q: what is the probability that this belief *b* is true, given that it was generated by "the BGP *P*"?
- $\Pr(b \mid b \text{ was generated by "the BGP } P")$
- This is supposed to be an *objective/physical reliability* probability, *not* a subjective or epistemic one that's relative to what we know/believe, *etc.*
- But, *which* objective probability is it?

The Analysis of Knowledge II Knowledge as Reliably Grounded True Belief IV

- Analogy: what is the probability that John has lung cancer (*j*), given that he was the product of "the" historical process (**H**), leading up to his present?
- $\Pr(j \mid \text{John was the product of "the" history } H)$?
- The problem here is that there are *many* ways of *describing* **H**, seemingly yielding *different values*.
- For instance, **H** is such that John was once a chain smoker. On *that* description of **H**, \Pr seems *high*.
- But, **H** is *also* such that John has long been an avid runner, with no history of cancer in his family.
- On *this* description, the probability seems *low(er)*.
- If we make our description of **H** *too* specific to John, then the probability will be 1 or 0 (*trivial*).

The Analysis of Knowledge II Knowledge as Reliably Grounded True Belief V

- We need a way of specifying "the BGP" in each given case, which yields "the probability" that the belief in question would be true, given its BGP.
- This way of specifying "the BGP" cannot appeal to concepts such as justification or knowledge.
- It's supposed to be a "purely objective" and "non-epistemically loaded" way of specifying a BGP for the purposes of determining its reliability.
- That seems quite difficult.
- Moreover, there is an old problem lurking here — the problem of *defeat* (the Zagzebski problem).
- This is related to the "how reliable" question, above. It's (again) a matter of *how independent* the reliability of "the BGP *P*" and the truth of *b* are.

The Analysis of Knowledge II
Knowledge as Reliably Grounded True Belief VI

- Think about the lottery paradox again.
- *One* description of “the” BGP yielding my belief *b* that my ticket will lose (or has lost) is as follows:
 - A process in which *chance* is crucial to the outcome
- On *that* description, the BGP sounds *unreliable*.
- But, we can *also truly* describe the lottery BGP as:
 - A process that yields true beliefs *almost always*
- On *that* description, the BGP sounds *reliable*.
- But, we are inclined to say that *b* is *not* known.
- As such, the first description must be more “apt”.
- But, what *criterion* for specifying descriptions could systematically ensure this choice?

The Analysis of Knowledge II
Knowledge as Reliably Grounded True Belief VII

- Note: “chance processes” *can* lead to *knowledge*.
- It was by chance (an astronomical coincidence) that I saw you on the train to Paris. But, my seeing *can* lead to knowing you were on the train.
- We need a way of saying *which* sorts of chance processes rule-out knowledge and which don’t.
- One way is to require that the agent *track the truth*.
- This “tracking” idea was first proposed by Nozick (*Philosophical Explanations*), and recently defended by our own Sherri Roush (*Tracking Truth*)
- The (rough) idea is to require *S* truly believe *p* **and**
 - If *b* hadn’t been true, *S* wouldn’t have believed *b*.
 - If *b* had been true, *S* would have believed *b*.

The Analysis of Knowledge II
Digression on Truth-Tracking I

- Note: truth-tracking accounts seem to avoid Gettier & lottery examples, *via* their second clause.
- Zagzebski’s doctor would have believed her patient had virus *V* *even if he hadn’t*.
- I would have believed that my lottery ticket was going to lose, *even if it had been the winner*
- If Zagzebski is right, then truth-tracking (**X**) must have a *very tight connection* with *b*’s truth. It does!
- All standard theories of counterfactuals (including Nozick’s) satisfy the following logical condition:

If $p \Box \rightarrow q$, then $p \supset q$.

- This (just as Zagzebski had already predicted) implies that **XB** entails T, where **X** is truth-tracking!

The Analysis of Knowledge II
Digression on Truth-Tracking II

- Here’s a proof that the second clause of truth-tracking (**X**) ensures that **XB** entails T.

1. Bb [(B): Assumption that *S* believes *b*]
2. $\sim b \Box \rightarrow \sim Bb$ [(X): Assumption of clause 2 of tracking]
3. If $p \Box \rightarrow q$, then $p \supset q$. [Logic of $p \Box \rightarrow q$]
4. If $\sim b \Box \rightarrow \sim Bb$, then $\sim b \supset \sim Bb$. [Instance of (3)]
5. If $\sim b \Box \rightarrow \sim Bb$, then $Bb \supset b$. [(4), contraposition for \supset]
6. $Bb \supset b$ [(2), (5), MP for “If ... then”]
7. b [(T): from (1), (6), MP for \supset]

- So, truth-tracking avoids Gettier-style examples, but only because it’s an **XTB** theory where **XB** entails T.
- Truth-tracking has another key consequence, as well.

The Analysis of Knowledge II Digression on Truth-Tracking III

- The truth-tracking theory *implies* that deductive inference does *not* transmit knowledge.
- One can track the truth of *b* without tracking various (obvious/known) logical consequences of *b*.
- Recall Dretske's horse/zebra example:
 - **Me:** Do you know that's a zebra?
 - **You:** Yes, of course.
 - **Me:** Do you know it's *not* a cleverly painted horse?
 - **You:** No, I guess not.
- In this example/context, you *track* that there's a zebra in front of you, but you *don't track* that there's *not* a cleverly painted horse in front of you.
- You'd believe *that even it were* a painted horse!

The Analysis of Knowledge II Digression on Truth-Tracking IV

- Our own Sherri Roush (*Tracking Truth*) defends an amended version of Nozickian truth-tracking.
- She adds a *closure principle* (to avoid Dretske by ensuring that deduction transmits knowledge):
 - If *S* knows that *p*, and *S* knows that *p* entails *q*, then *S* also knows that *q*.
- And, she moves away from counterfactual conditionals to *conditional probability thresholds*:
 - $\Pr(Bp \mid \sim p)$ is low (threshold determined by context)
 - $\Pr(Bp \mid p)$ is high (threshold determined by context)
- This makes her theory susceptible to Gettier/Zagzebski cases, but it avoids Nozick's *infallibilism*.
- And, Roush now faces a *generality/specification* problem, if the probabilities are *objective*.

The Analysis of Knowledge II Knowledge as Reliably Grounded True Belief VIII

- This suggests a new dilemma for theories of K:
 - Either your theory is *fallibilist*, in which case it falls prey to Gettier/Zagzebski examples, or it's *infallibilist*, in which case it falls prey Dretske-style violations of deductive K-transmission
- Enter "relevant alternatives" theory, which gives a way to respond to K-transmission failure cases.
- In the Dretske zebra example, a "relevant alternative" (that it's a cleverly painted horse) is present, but *you cannot discriminate it from b*.
- Relevant alternatives theory requires that you be able to discriminate *b* from all relevant alternatives
- This would help with the zebra example (and also the Jane's twin example), and it also gives a way to respond to *skepticism* (we'll return to *that* later)

The Analysis of Knowledge II Knowledge as Reliably Grounded True Belief IX

- But, what *counts as* a "relevant alternative"? And, what's involved in *discriminating* alternatives?
- Here, some have suggested that which alternatives count as "relevant" will depend on the *context*.
- On the *contextualist* view, in *high standards* contexts, more things will count as relevant alternatives than in *low standards* contexts.
- Philosophy (historically) has presupposed some dazzlingly *high standards* for knowledge.
- The fact that we're inclined to say that the doctor in Zagzebski's situation *does not know* the patient has virus *V* shows how high these standards are.
- Contextualism also gives us a nice explanation of what's going on in the Dretske zebra dialogue

The Analysis of Knowledge II Knowledge as Reliably Grounded True Belief X

- The Contextualist/relevant alternatives explanation of the Dretske zebra/horse dialogue is as follows:
 - At the beginning of the dialogue, the standards are not so high. So, the alternative of “not cleverly painted horse” is *not* relevant.
 - At the end of the dialogue, the standards have shifted (they’re much higher), and the alternative “not cleverly painted horse” becomes *relevant*.
- In the dialogue, I’m the philosopher. *Annoying!*
- As we’ll see later, this also gives an explanation of what happens in *skeptical* contexts. There, the standards are *so* high as to render K *impossible*.
- We should also think about what is required for *discrimination* of alternatives. This must involve *more than tracking*. But, what could that be?

The Analysis of Knowledge II Prelude to Internalism vs Externalism I

- Until today, we’ve (usually) been presupposing that justification is *a* component of knowledge.
- This reflects a historical *internalist* bias.
- Recall that justification (as we have been conceiving it) has an *internal* component.
- We use “justified” in a kind of “by your own lights” way. But, we use “knows” in more *external* way.
- What if I (in fact) *do track the truth* about *b*, and I believe *b*, and *b* is true (and there is no “bad luck” involved, *etc.*). Must I be *justified in believing b*?
- What if I (falsely) believe that I *do not track the truth* about *b* (say, on testimony from a reliable psychologist). Presumably, I would *not be justified in believing b* (rather, in *not-b!*). But, on a (pure) *truth-tracking* account, I nonetheless can *know b*.

The Analysis of Knowledge II Prelude to Internalism vs Externalism II

- The naturalistic (*e.g.*, truth-tracking and reliabilist) accounts of knowledge tend to be *externalist*.
- But, the JXTB theories tend to be *internalist*.
- An externalist will just require true belief & something *external* to the agent, like *reliability*.
- An internalist will require true belief & *justification*, with an *internal* “by your own lights” component — something which is “internally accessible” to S.
- One is either reliable/truth-tracking or not. One needn’t have “internal access” to *any* information about *whether* one is reliable or truth-tracking.
- The externalist *just* requires the right sort of objective/external connection to the truth of *b*, *without any* “by your own lights” requirements.
- Question: is “discrimination” internal or external?

Second Essay Questions I

1. Compare and contrast classical vs empiricist views of truths of reason (with any eye toward deciding which view you think is more plausible). Your discussion should begin by addressing the central question: (a) What are truths of reason, and why are they necessarily true on a classical view? [Here, be sure to cover the synthetic a priori as well as the analytic]. Additionally, you should also address at least one of the following three questions (from the point of view of classicism vs empiricism): (b) From an empiricist standpoint, why is the synthetic a priori an illegitimate category?, (c) Can there be empirical necessary truths?, and (d) Are all conceptually necessary truths a priori?
2. Compare & contrast Coherentism & Foundationalism (with any eye toward deciding which view you think is more plausible). Your discussion should touch on *at least* the following two issues: (a) the precise role of coherence as a contributor to justification (and/or truth and knowledge), (b) whether foundational beliefs (or the coherence of one’s beliefs) are (is) the *sine qua non* of justification (and/or knowledge).

Second Essay Questions II

3. Discuss either the preface paradox or the lottery paradox (but not both). First, explain what the (chosen) paradox is. Then, explain what it is supposed to show about deductive/inductive inference and the transmission of justification/knowledge. Do you think it succeeds in showing what it is supposed to show? Explain. [See lectures #12 and #13 for my notes on these two paradoxes.]
4. Deductive inference is supposed to be a primary means of expanding our knowledge. Question this, by asking whether deductive inference (even) transmits knowledge. Describe and explain example(s) that are supposed to show that deductive inference does not (even) transmit knowledge. What do you think about such examples? Can you think of a response on behalf of someone who thinks deductive inference does/should transmit knowledge? How serious a problem would it be if knowledge were not (in general) transmitted by deductive inference?

Second Essay Questions III

5. Discuss some Gettier-style counterexample(s) against the justified true belief (JTB) theory of knowledge. Here, you can discuss Gettier's original example(s) (I discuss these two examples in lecture #16), or you can discuss other similar example(s) that you think are even more compelling. You should give a precise analysis and discussion of your chosen example(s), and you should say whether you think they are effective (for refuting the JTB theory). Can you think of a response on behalf of a JTB theorist? In light of such examples, what do you think about the prospects for a JTB theory of knowledge (here, the Zagzebski paper might be especially useful — see lecture #16 for more on that paper, and [here](#) is a JSTOR link to the paper itself)?