

# Announcements and Such

- One Song — *Mark Knopfler*
  - “Why Aye Man” from *The Ragpicker’s Dream*
- First essays to be returned today (after class)
- Second essay will be assigned on March 20th.
- Today: Part I of three parts on the Architecture of Knowledge (*very* serious theoretical epistemology)
- First, one leftover from last time...

## Inference and the Extension of Knowledge Inferential Transmission of Knowledge Revisited I

- It seems that knowledge *can* be transmitted, *even through inductive inference*, and even if some justification/probability is “lost in the inference”.
- Last time, I used one of Audi’s examples to illustrate this point, which wasn’t so compelling.
- But — so long as we have *some* knowledge that has been transmitted (at *some* stage) through *some* inductive (scientific) inference — the point stands.
- The alternative is *inductive skepticism* (more later).
- Moreover, knowledge can *fail* to be transmitted even when the probabilities are *very high*.
- You *know* that you hold one of a million tickets in a fair lottery, which will have *one* winner. You infer — *with very high probability* — 0.999999, that you will lose. You do not *know* you will lose.

## Inference and the Extension of Knowledge Inferential Transmission of Knowledge Revisited II

- Here’s the *reductio* argument for the claim that you do not know that you won’t win the lottery:
  1. Assume that you *do* know that you won’t win the lottery. [This is our *reductio* assumption.]
  2. Then, *each* ticket holder knows that they won’t win the lottery. [Because there’s nothing special or privileged about *your* epistemic situation.]
  3. Since knowledge implies truth, (2) implies:  
ticket #1 won’t win &  
ticket #2 won’t win &  
& ... &  
ticket #1,000,000 won’t win.
  4. But, (3) implies that *no ticket will win*, which *contradicts* the set-up of the example. *QED*.
- A few further remarks are in order here.

## Inference and the Extension of Knowledge Inferential Transmission of Knowledge Revisited III

- We’re assuming that *knowledge implies true belief*:
  - *S* knows that  $p \Rightarrow S$  believes that  $p$ , and  $p$  is true
- But, we’re *not* assuming that knowledge is *infallible*. That is, we’re allowing for:
  - *S* knows that  $p \not\Rightarrow S$  has *conclusive grounds* for  $p$
- We’re also *not* assuming that knowers have *second-order* knowledge. That is, we’re allowing for:
  - *S* knows that  $p \not\Rightarrow S$  knows that *S* knows that  $p$
- This is the denial of what is called the KK-principle. KK plays a crucial role in arguments for skepticism. We’ll discuss it at length later on.
- OK, moving on to the architecture of knowledge...

### The Architecture of Knowledge I Foundationalism – The Classical Position

- This chapter is concerned with how an epistemic agent *S*'s knowledge (and beliefs) is *structured*.
- *Foundationalism* claims that (*e.g.*) *S*'s perceptual knowledge forms the *basis/foundation* for *all* of *S*'s knowledge (with the rest being *indirect*).
- We will examine foundationalism and some of the alternatives to it that have been floated.
- Before we do that, we will think about the nature of inferential chains and the structure of belief.
- Intuitively, some of our beliefs are *direct* (*e.g.*, certain simple perceptual beliefs), and others are *indirect* (*i.e.*, based on other beliefs).
- This naturally leads to a (weak) foundationalism, according to which *S*'s *direct* beliefs/knowledge form a foundation for *all* of *S*'s bel./knowledge

### The Architecture of Knowledge I Infinite Inferential Chains? I


- One way that foundationalism could fail is if *infinite* inferential chains were possible.
- It is very difficult to see how this could be the case (for any actual *S*), for various reasons.
- First, *S* cannot *believe* things *at will*, nor can *S* decide what her beliefs are *based on at will*.
  - Some beliefs seem *forced* upon us (*e.g.*, certain perceptual beliefs, we surely don't *choose*)
  - If something seems highly implausible to you, you can't always "will yourself to believe it"
  - If your grounds for believing *p* are perceptual, you can't "will them to be testimonial", *etc.*
- That is, *unrestricted doxastic voluntarism* is *false*.
- There are other problems with infinite chains.

### The Architecture of Knowledge I Infinite Inferential Chains? II

- Can *S* even *have* infinitely many beliefs to serve as elements of such an infinite inferential chain?
- I don't see how. But, even if there are some domains where infinitely many propositions are around (logic/math), it's hard to see how they could form an infinite *inferential chain* for *S*.
- What would such a chain look like?
- Audi considers an alternative to foundationalism in which each belief is based on a proposition concerning how things "seem" to *S*.
- Perhaps *S* could have infinitely many beliefs like this (...it seems to *S* that it seems to *S* that it seems to *S*...). But, the trouble is that each one would have to be *based on* last one in the chain.
- This gets rather absurd rather quickly...

### The Architecture of Knowledge I Circular Inferential Chains? I

- Another way foundationalism could fail is if *circular* inferential chains were possible.
- Then, *S* could have finitely many beliefs, but each one depending on another one (thus all *indirect*).
- Example: consider the following "circular chain"
  - I believe (*p*) there is a wind. I believe *p* on the basis of my believing (*q*) there is a swaying of the trees; I believe *q*, on the basis of my believing (*r*) I have an impression of such swaying; and I believe *r*, on the basis of believing *p* (the first belief).



$y \longrightarrow x$   
*S* believes *x* on  
 the basis of *y*

### The Architecture of Knowledge I Circular Inferential Chains? II

- Such circular chains are incompatible with (a)–(c):
- (a) *S*'s belief that *p* is based on *S*'s belief that *q* only if *S*'s believing that *q* is (at least in part) causally responsible for *S*'s believing that *p*.
- *E.g.*, if I believe there is a wind, on the basis of my believing that the trees are swaying, then I believe that there is a wind (at least in part) because I believe that the trees are swaying.
- (b) If *x* is (in part) causally responsible for *y*, and *y* is (in part) causally responsible for *z*, then *x* is (in part) causally responsible for *z*.
- (c) nothing is (even in part) causally responsible for *itself*.
- Condition (a) is a theoretical assumption/definition of this course. But, (b) & (c) are more controversial.

### The Architecture of Knowledge I Circular Inferential Chains? III

- (b) says (more or less) that *causation is transitive*: if *x* causes *y* and *y* causes *z*, then *x* causes *z*.
- Is this right? Consider the following example:
  - A boulder is dislodged, and begins rolling ominously toward Hiker. Before it reaches him, Hiker sees the boulder and ducks. The boulder sails harmlessly over his head with nary a centimeter to spare. Hiker survives his ordeal.
  - Here, the falling of the boulder (*x*) is (in part) causally responsible for (*y*) Hiker ducking, which is (in part) causally responsible for (*z*) Hiker surviving the ordeal. But, it seems odd to say that the falling of the boulder is (*even in part*) causally responsible for Hiker's survival.
- There is a lot of interesting literature on this question (Ned Hall and Chris Hitchcock).

### The Architecture of Knowledge I Circular Inferential Chains? IV

- Assumption (c) could also be questioned, but it is far less controversial than assumption (b).
- In any case, let's just grant (b) and (c), *arguendo*.
- Then, circular chains are ruled-out, since if *p* is based on *q*, which is based on *r*, which is based on *p*, then (a) & (b)  $\Rightarrow$  *x* is causally responsible for *x* (where *x* = *S*'s believing *p*), which *contradicts* (c).
- Surely, it would be strange to think of beliefs "holding themselves up by their own bootstraps".
- But, it seems to me that the proponent of circular inferential chains could simply *deny* (b).
- That said, it is helpful to be more precise about what the basing relation really involves.
- 3 components: conviction, explanation, memory.

### The Architecture of Knowledge I Circular Inferential Chains? V

- If *S* believes *p* on the basis of *q*, then (often):
  - *S*'s belief in *q* tends to increase or buttress *S*'s conviction in *S*'s belief that *p*, and
  - *S*'s belief in *q* tends to *explain* (in part and in some broadly causal way) *why S*'s believes *p*, and
  - *S*'s belief in *q* tends to help *S* remember *p*.
- This is not the kind of support relation that a belief may be plausibly thought to bear *to itself*.
- Moreover, somewhat *independently* of the causal responsibility argument — *via* (a) and (b) — it seems strange to say that a *p,q,r* chain like the one above could satisfy all three roles of "basing".
- So, circular inferential/belief chains do seem somewhat odd in the end, I think.

### The Architecture of Knowledge I The *Epistemic* Regress Problem I

- What we really care about are possible *epistemic* dependencies among beliefs (and knowledge).
- We will focus on the structure of *knowledge*.
- Could all of our *knowledge* be *indirect* — that is, could it all be based on *other* knowledge we have?
- There are four (standard) kinds of *epistemic chains*:
  - *Infinite* epistemic chains (knowledge that *p* is based on *K/B* that *r*, which is based on *K/B* that *q*, ...)
  - *Circular* epistemic chains
  - Epistemic chains terminating in *non-knowledge*
  - Epistemic chains terminating in *knowledge*
- Indeed, it is usually assumed that these are the *only* possibilities, which leads to an *epistemic regress argument*. We'll look at each, then the argument.

### The Architecture of Knowledge I The *Epistemic* Regress Problem II

- Infinite epistemic chains are even more absurd than infinite inferential chains.
- It's just *bizarre* to claim that *knowledge* ever emerges from an (*essentially*) infinite chain.
- What about *circular epistemic* chains? How could knowledge emerge from a *circular* epistemic chain?
  - I know (*p*) there is a wind. I know this on the basis of (*q*) the trees are swaying. Now I *think* I know *q* because I *see* them swaying. But, perhaps my seeing this is only the *causal* basis of my belief that *q*, and I just do not notice that it is only on the basis of, say, my knowledge that (*r*) I have a visual impression of swaying that I know *q*. And, perhaps I know *r* on the basis of *p*, which is the first knowledge in the chain.
- But, *how* does my knowing *p* *epistemically* ground my belief that *r*? *That* seems like an *introspective* belief, *not* an inferential one, based on *p*.
- These seem like odd possibilities — at best.

### The Architecture of Knowledge I The *Epistemic* Regress Problem III

- What about an epistemic chain that terminates with *non-knowledge*. Is this possible?
- Clearly, *some* such candidates are not on:
  - Suppose that (in foggy conditions) I simply *guess* that what I see is a swaying of trees, but happen to be right. Might I then *know* there is a wind anyway, provided there is? Surely not; knowledge cannot be grounded in *mere guesswork* — *even when the guess is correct*.
- OK, but what if I make an *educated* guess?
- In such a case, perhaps we *could* have knowledge emerge (assuming we're not inductive skeptics!).
- But, when that happens, the chain will terminate in something *pretty close to knowledge* (*firm sand*!).
- That is, we'll have some *substantial degree of justification* in the terminal node of the chain.

### The Architecture of Knowledge I The *Epistemic* Regress Problem IV

- Finally, epistemic chains that terminate in *direct knowledge* (or *foundational* knowledge).
- Here, the anchor of the chain arises, *directly*, from perception, memory, introspection, or reason.
- Such chains can vary along four dimensions:
  - *Composition*: the *sorts* of beliefs constituting them.
  - *Transmission Types*: the links may be deductive, inductive, or combine both kinds of links.
  - *Ultimate Grounds* (grounds of the anchors): they may be grounded experientially or rationally.
  - *Justificational Strength*: the degree of justification they give to the belief (at the other end).
- Note: the anchors *need not be self-evident* or *indefeasibly justified* (as in "old-time foundationalism").

### The Architecture of Knowledge I The Epistemic Regress Argument I

- Now, we're ready for the epistemic regress argument, which can be seen as an argument for foundationalism (at least of a weak kind):
  1. If *S* knows *p*, then *p* occurs in an *epistemic chain* (all knowledge occurs in *some* epistemic chain — with direct knowledge being the *basis case*).
  2. The only possible kinds of epistemic chain are the four mutually exclusive kinds just discussed: the infinite, the circular, those terminating in beliefs that are not knowledge, and those terminating in direct knowledge.
  3. Knowledge can emerge only in the 4th kind.
  4. If *S* has *any* knowledge, then *S* has *some* direct knowledge (if there is any knowledge, there is some direct knowledge).

### The Architecture of Knowledge I The Epistemic Regress Argument II

- Aristotle offered a similar argument, which has since been repeated/copied many times.
- If the epistemic regress argument works, it shows something *stronger*: that all knowledge is *traceable back to* some direct knowledge as its foundation.
  - A similar argument could be run for *justification* instead of knowledge: if *S* has any *justified* beliefs, then *S* has *some directly justified* beliefs.
- Let's just grant premises (1) and (2).
- We have already seen reasons to worry about premise (3). But, *even if* (3) is false, the *justification* version of the argument still seems to go through, with (3) replaced by something a bit weaker:
  - (3\*) justified belief can only emerge in *either* the 3rd *or* the 4th kinds of epistemic chains.

### The Architecture of Knowledge I Foundationalism & Coherentism I

- Various kinds of foundationalism are compatible with being motivated by this sort of argument.
  - *Strong* foundationalism: indirectly justified beliefs get *all* their justification from F.
  - *Moderate* foundationalism: indirectly justified beliefs would not have any justification were it not for F, but they can get *some* of their justification from *other* sources (not in chain).
- The epistemic regress argument also leaves open the *content* or *kind* of propositions comprising F.
- The main alternative to foundationalism is *coherentism*, which says (broadly construed):
  - The justifiedness of a belief depends on its *coherence with other beliefs one holds*.
- We'll discuss coherentism next time...