Phil. 256: Becausal Relations
Dec. 8, 2021

Meeting 11: An Application in Epistemology

I. The Causal Theory of Knowledge

Goldman's causal theory of knowledge (from his 1967 article) has had an odd legacy:

> On the one hand, Goldman himself disowned the theory in 1976, and causal theories of knowledge have virtually no defenders in contemporary epistemology.
>
> On the other hand, Goldman's proposal has had an outsized influence on fields other than epistemology: standard arguments against mathematical Platonism (in philosophy of mathematics) and normative non-naturalism (in metaethics) often appeal to some sort of causal theory of knowledge, whether directly or indirectly (e.g. when it is assumed that only causal connections can make it the case that a belief has the factors—such as safety or sensitivity—required for knowledge).

Goldman's proposal:

> *the causal theory of knowledge*: S (empirically) knows that *p* iff S's belief that *p* is causally connected in an "appropriate" way with the fact that *p*.
>
> > Note: Goldman only intends this to be an analysis of empirical knowledge: "I think that the traditional analysis is adequate for knowledge of nonempirical truths" (p. 357).

Exactly which causal connections count as appropriate? Goldman's basic methodology involves pointing to certain paradigm cases and then leaving it to science to fill in the details of the relevant causal processes.

An example of a causal connection that *does not* count as appropriate:

- *Blow to the Head*: I fall down the stairs and hit my head. The blow to my head scrambles my brain, causing me to form all sorts of wild beliefs, including a belief that I have just hit my head. [I have just hit my head] causes [I believe that I have just hit my head], but not in an appropriate way.

> Note: for simplicity's sake, we will be assuming that facts can be causes and effects.

Some examples of causal connections that *do* count as appropriate:

- *Perception*: I see a wall in front of me. As a result, I come to believe that there is a wall in front of me. In this case, [There is a wall in front of me] causes [I believe that there is a wall in front of me] in an appropriate way.

- *Inference (Pattern 1)*: ". . . S perceives that there is solidified lava in various parts of the countryside. On the basis of this belief, plus various 'background' beliefs about the production of lava, S concludes that a nearby mountain erupted many centuries ago" (p. 361).

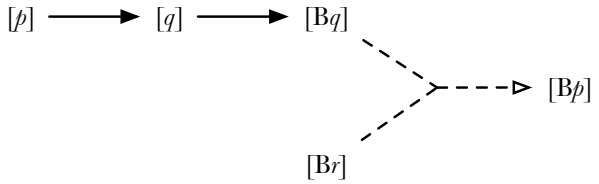

> [*p*] = the fact that a nearby mountain erupted many centuries ago
> [*q*] = the fact that there is solidified lava throughout the countryside
> [*r*] = a "background" fact about "the ways in which lava is produced and how it solidifies"
> [B*p*] = the fact that S believes that *p*
>
> The causal chain from [*p*] to [B*p*] counts as appropriate because S has made a series of inferences that properly mirror the causal chain leading to S's initial belief.

> When a structure of beliefs in a subject's head mirror a causal chain in the world, the subject is said to have "reconstructed" the causal chain. The subject's reconstruction must "contain no mistakes," and though the subject "need not reconstruct *every* detail of the causal chain, he must reconstruct all the important links" (p. 363).
>
> Also, "the knower's inferences must be warranted[:] the propositions on which he bases his belief of *p* must genuinely confirm *p* very highly, whether deductively or inductively" (ibid.).

- *Inference (Pattern 2)*: S sees a fire in the fireplace and, as a result, comes to believe that there is smoke coming out of the chimney.

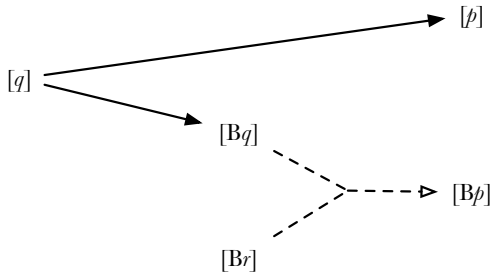

> [*p*] = the fact that there is smoke coming out of the chimney
> [*q*] = the fact that there is a fire in the fireplace
>
> [*q*] is a common cause of [*p*] and [B*q*], the latter of which leads (via inference) to [B*p*]. Such a structure of beliefs in the head is another way of properly mirroring a causal structure in the world.

"Appropriate," knowledge-producing causal processes include (but are not necessarily limited to) perception, Pattern 1 and 2 inferences, and any combination of perception and Pattern 1 and 2 inferences.

Note that Goldman has ready responses to three common objections to causal theories of knowledge:

- *objection #1*: What about deviant causal chains? *solution*: These chains will feature important links that the subject has not reconstructed.
- *objection #2*: What about knowledge of the future? *solution*: Such knowledge fits Pattern 2.
- *objection #3*: But isn't allowing common causes too permissive? Presumably the Big Bang is a common cause of almost every contingent fact and corresponding belief. *solution*: Most subjects have not reconstructed even a tiny fraction of the important links from the Big Bang to their belief.

Some aspects of Goldman's proposal that need greater clarification and/or require some reformulation:

- What exactly happens at the initial mind–world nexus? In the case of visual perception, does it need to be that [*p*] causes [It visually seems to S that *p*], which in turn causes [S believes that *p*]?
- Is it enough for each link in a knowledge-producing causal chain to involve partial causation, or does the subject need to reconstruct all (or most) of the partial causes that together constitute the full cause of a given effect in that chain?
- What exactly is the content of a given "background belief"?
- Doesn't a background belief itself need to constitute knowledge in order for the belief inferred from it to constitute knowledge?

  But if it does need to be, then (since every piece of inferential knowledge relies upon at least one background proposition) either (a) there are some background facts that one can be in causal contact with via perception (or some similar causal mechanism), or (b) there are ways of coming to know background facts not covered by the causal theory, or (c) inferential knowledge is impossible.

II. A Problem

A serious problem for the causal theory of knowledge arises when we consider our knowledge of existential and universal generalizations. Presumably I know the following facts: [Someone in this room owns a dog] and [All humans are mortal]. But neither fact seems causally connected with my beliefs.

In response, Goldman proposes that we allow the structure of facts in the world (which the subject must mirror via her beliefs) to include *either* causal *or* logical relations among the nodes in that structure.

> He justifies this move by appealing to the following principle (p. 368; 'GP' = 'Goldman's Principle'):
>
> GP. If *x* is logically related to *y*, and if *y* is a cause of *z*, then *x* is a cause of *z*.

Let us call the theory after it has been modified in this way *the logico-causal theory of knowledge*, since now it's no longer a pure causal theory. Goldman claims that, with this modification in place, the following structures suffice for knowledge of [Someone in this room owns a dog] and [All humans are mortal]:

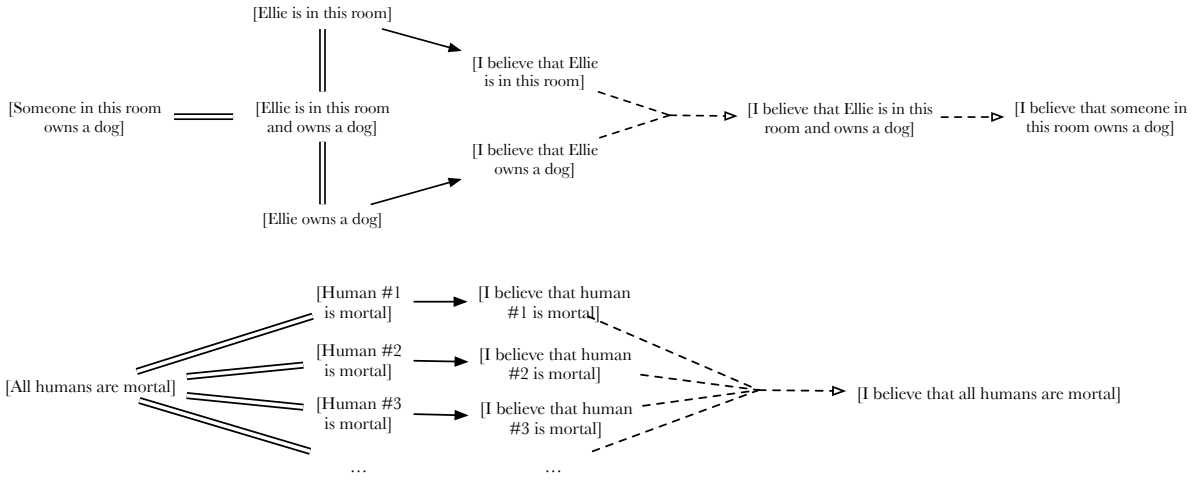

But for this emendation of his theory to work, sometimes Goldman needs 'is logically related to' to mean 'entails', and other times he needs it to mean 'is entailed by'. And in neither case is (GP) remotely plausible.

> GP*. If *x* entails *y*, and if *y* is a cause of *z*, then *x* is a cause of *z*.

Counterexample: let *x* = [I dropped the chalk, and there are infinitely many primes], *y* = [I dropped the chalk], and *z* = [The chalk broke].

> GP**. If *x* is entailed by *y*, and if *y* is a cause of *z*, then *x* is a cause of *z*.

Counterexample: let *x* = [{Chalky} exists], *y* = [I dropped Chalky], and *z* = [Chalky broke].

Moreover, for each of these counterexamples to (GP), we can construct an analogous counterexample to the logico-causal theory of knowledge.

> Surely I can't come to know [I dropped the chalk, and there are infinitely many primes] by inferring it from [I dropped the chalk].
>
> Surely I can't come to know [{Chalky} exists] by inferring it from [I dropped Chalky]. (Or, at least, surely a broadly causal theory of knowledge shouldn't allow this sort of an easy ontological inference.)

So we're in a bind: we need to account for knowledge of existential and universal generalizations, but allowing logical relations in addition to causal relations is too coarse-grained a way of doing so.

III. A Solution

The obvious fix: let's appeal to grounding instead of entailment.

Then we can appeal to two sorts of becausal relations in our structures in the world: the 'because' of causation and the 'because' of grounding.

> Here analogues of Goldman's principle are very plausible, I claim. ('BP' = 'Becausal Principle'.)
>
> BP_1. If *p* because$_{grounding}$ *q*, and *q* because$_{causation}$ *r*, then *p* because *r*.
>
> E.g. let [*p*]= [Xanthippe is a widow], [*q*] = [Socrates died], and [*r*] = [Socrates drank the hemlock].
>
> BP_2. If *p* because$_{causation}$ *q*, and *q* because$_{grounding}$ *r*, then *p* because *r*.
>
> E.g. let [*p*] = [The chalk broke], [*q*] = [I dropped the chalk], and for [*r*] insert your favorite theory of that in virtue of which I count as performing my action of dropping the chalk.

And rather than appealing to an inferential relation among our subject's beliefs, we can instead appeal to the 'because' of basing, which includes inferential basing but also more.

> So in our graphs from now on, let a solid arrow represent the 'because' of causation, a dotted arrow represent the 'because' of grounding, and a dashed arrow represent the 'because' of basing.

Now the explanation of how I know [Someone in this room owns a dog] is straightforward: it is a case of Pattern 2 inferential knowledge, with both grounding and causal links among the facts in the world.

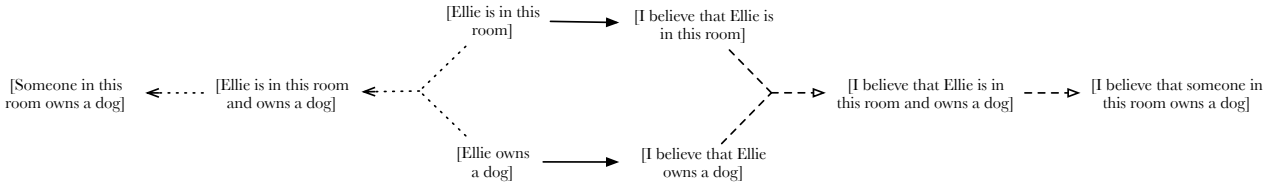

When we make all of these changes, the causal theory of knowledge has now become

> *the becausal theory of knowledge*: S knows that *p* iff S's belief that *p* is becausally connected in an appropriate way with the fact that *p*,

where an appropriate "becausal chain" between the fact that *p* and a belief that *p* is allowed to have either grounding or causal links on the worldly side of things and has basing links on the mental side of things, and where the mental nodes mirror the worldly nodes in either a Pattern 1 or a Pattern 2 manner.

Some virtues of the becausal theory:

1. *Unlike Goldman's logico-causal theory, the becausal theory is well motivated.*

   The basic idea behind Goldman's original, pure causal theory was that, in order to have knowledge, the dependency structures in one's mind must mirror the dependency structures in the world; but why should those worldly dependencies be restricted to causation?

2. *Unlike Goldman's pure causal theory, the becausal theory readily accounts for knowledge of existential and universal generalizations, of disjunctions and conjunctions, and more.*

   On the becausal theory, one can come to know all of the following by basing one's belief in that fact on one's belief in the grounds of that fact: [There is a seminar is taking place in this room], [Object *x* has determinable property *d*]; [It was wrong of person *y* to perform action *a*]; [*c* caused *e*].

3. *There is no need to restrict the becausal theory so that it is only an account of empirical knowledge.*

   On an *empiricist* version of the becausal theory, *perception* is the only allowable mind–world nexus, and all other knowledge is erected on top of that which one directly knows via perception.

   On a *rationalist* version of the becausal theory, non-perceptual forms of mind–world nexus such as *intuition* are allowed.

More generally, I claim that the becausal theory of knowledge solves some of the major problems besetting the causal theory while introducing no new problems.

- *potential new problem #1*: What about Pattern 1 inferences featuring both causal and grounding links? Isn't it problematic to say that one can come to know [p] by inferring it from [$p \vee q$] (on its own), or come to know [The Harvard insignia is crimson] by inferring it from [The Harvard insignia is red]?

  an unpromising response: Deny that disjunctive facts and determinable color facts have causal powers (or, at least, the sorts of causal powers that can underwrite a perceptual mind–world nexus).

  a better response: Just as all inferences in the causal theory are mediated by a (known) background theory that, together with the inference's premises, highly confirms the inference's conclusion, so too must inferences in the becausal theory be mediated by a (known) background theory that, together with the inference's premises, highly confirms the inference's conclusion. And in the problematic inferences in question, there's no such background theory (unless, for instance, [$p \vee q$] = [$p \vee p$], in which case the inference seems fine).

- *potential new problem #2*: But isn't logical inference from known premises always a way of coming to know the conclusion? Not all such inferences are cases in which the conclusion is grounded in the premises.

  reply: First, some epistemologists do in fact deny that performing a logical inference always extends one's knowledge.

  Second, some logical inferences will be cases in which one of the premises is grounded in the conclusion, as in the inference from [$p \vee q$] and [$\sim p$] to [q], and so (if mediated by the proper background theory) will count as a Pattern 1 inference.

  Third, some logical inferences (esp. more complex ones) yield a belief in the conclusion that is not directly based on the subject's beliefs in the premises, but rather features an indirect basing structure corresponding to the intermediate steps in the inference (as appreciated by the inferrer). This greater structure on the mental side of things allows for greater structure in the possible grounding relations between premises and conclusions on the worldly side of things that might make for knowledge.

    For instance, if I infer [The flag is red] from [The flag is crimson or cyan] and [The flag isn't cyan], presumably it's not the case that my belief in <The flag is red> is directly based on my belief in <The flag is crimson or cyan> and my belief in <The flag isn't cyan>, but rather stands in an indirect basing relation that is mediated by a belief in <The flag is crimson>.

  Putting these three points together, the hope is that all genuinely knowledge-extending logical inferences can be accounted for by means of direct or indirect grounding paths between the premises and the conclusion.

    a challenging case: Suppose I infer [The flag is red or blue] from [The flag is crimson or cyan] without believing any of the disjuncts.

    reply #1: Insist that [The flag is crimson or cyan] itself grounds [The flag is red or blue], even according to a Fine-style impure logic of grounding they instead have a common ground.

    reply #2: Add suppositional mental states into our basing structure, and then insist when I perform this inference, *a supposition that the flag is crimson* and *a supposition that the flag is cyan* are jointly based on *a belief that the flag is crimson or cyan*, which in turn jointly lead to *a supposition that the flag is red* and *a supposition that the flag is blue*, which in turn jointly lead to *a belief that the flag is red or blue*.

    an even more challenging case (due to Fine): Suppose I infer [$q \vee r$] from [$(p \,\&\, q) \vee (r \,\&\, s)$] without believing <$q$> or <$r$>. (Here a version of the suppositional strategy seems the only option.)

IV. An Upshot

An important upshot of the becausal theory of knowledge: *it is directly poised to address epistemological challenges to normative non-naturalism*.

Normative non-naturalists hold that there are causally-inert normative facts that are different in kind from the natural facts. So how do non-naturalists (who are not normative skeptics) explain how it is that our normative beliefs tend to correlate with the normative facts, or how it is that our faculties for making normative judgments are able to track the truth (despite being saturated by evolutionary influences)?

The standard reply is to offer a "third-factor explanation": there is some natural property that acts can have (for Nozick [1981]: *being helpful and not harmful to others*; for Enoch [2010]: *promoting one's own survival*) which both *causes* us to believe that the act has some normative property and *grounds* the fact that it is has this property.

On the becausal theory, the dependency structure advocated by third-factor theorists is exactly the dependency structure that can underwrite normative knowledge via a Pattern 2 inference.

- *objection #1*: If normative non-naturalism is true, then a subject cannot come to know the relevant normative principles that might serve as a background theory for a third-factor explanation.

  reply: Why can't a subject come to know the normative principles in question? Recall that even in the non-normative case, the causal (and becausal) theory faces problems accounting for a subject's knowledge of the initial background theories used to gain a first foothold in inferential knowledge. It's not clear why whatever is said to address that problem there can't also help here.

  > In both the normative and the non-normative case, we sometimes make use of background theories about entities, facts, and relations that are not themselves perceptually observable.
  >
  > In both the normative and the non-normative case, purely analytic background theories don't seem to be enough to get the entire enterprise of inferential knowledge started.

- *objection #2*: Third-factor explanations can at most account for our knowledge of normative facts that are (at least partially) grounded in the natural facts. So how do normative non-naturalists account for our knowledge of *the ungrounded normative facts*, to which they are also committed?

  reply: For rationalist becausal theorists, such knowledge is possible via intuition.

  For empiricist becausal theorists, such knowledge is possible *either* via the way we come to know our initial background theories *or* via a Pattern 3 inference: after a Pattern 2 inference to a partially-naturally-grounded normative fact (call it $[n_1]$), the dependency structure one is mirroring in one's mind zags down again to the ungrounded normative fact (call it $[n_2]$) partially in virtue of which that partially-naturally-grounded normative fact obtains.

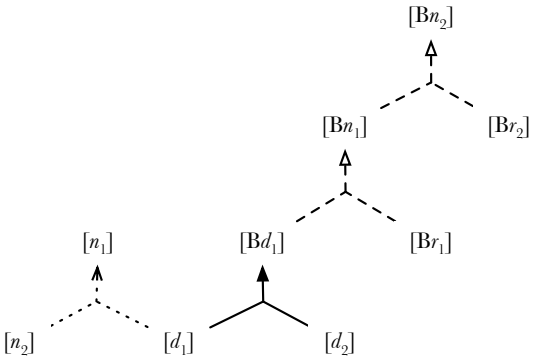

  More generally, if it is possible, according to the becausal theory, for us to know *ungrounded, uncaused non-normative facts* (such as facts about the Big Bang or facts about the fundamental laws of nature), then it is also possible for us to know *ungrounded, uncaused normative facts*.