

## Motivating the Problem

### 1. More on Lewis on Causation

If distinct events  $c$  and  $e$  both occur, then  $e$  *counterfactually depends* on  $c$  iff the following is the case: if  $c$  hadn't occurred, then  $e$  wouldn't have occurred.

*Lewis' original analysis of causation:* Event  $c$  is a cause of event  $e$  iff:

1.  $c$  and  $e$  both occur,
2.  $c$  and  $e$  are distinct, and
3. there is a (possibly empty) set of distinct events  $\{d_1, d_2, \dots, d_n\}$  such that  $d_1$  counterfactually depends on  $c$ ,  $d_2$  counterfactually depends on  $d_1$ ,  $\dots$ , and  $e$  counterfactually depends on  $d_n$ .

Potential counterexamples (from "Causation by Influence"):

- *late cutting:* Billy and Suzy both throw a rock at a bottle. Each throw would be sufficient on its own to break the bottle. However, Suzy's rock gets there just moments before Billy's does.
- *trumping:* "The sergeant and the major are shouting orders at the soldiers. The soldiers know that in case of conflict, they must obey the superior officer. As it happens, there is no conflict. Sergeant and major simultaneously shout 'Advance!'; the soldiers hear them both; the soldiers advance" (p. 183).

Event  $e$  is *fragile* if, or to the extent that, it could not have occurred at a different time, or in a different manner.

An *alteration* of event  $e$  is either a very fragile version of  $e$ , or else a very fragile alternative event that is similar to, but numerically distinct from,  $e$ .

If  $c$  and  $e$  are distinct actual events, then  $c$  *influences*  $e$  iff there is a substantial range  $c_1, c_2, \dots$  of different not-too-distant alterations of  $c$  (including the actual alteration of  $c$ ) and there is a substantial range  $e_1, e_2, \dots$  of not-too-distant alterations of  $e$  (at least some of which differ) such that the following holds: if  $c_1$  had occurred, then  $e_1$  would have occurred; if  $c_2$  had occurred, then  $e_2$  would have occurred; and so on.

*Lewis' emended analysis of causation:* Event  $c$  is a cause of event  $e$  iff:

1.  $c$  and  $e$  both occur,
2.  $c$  and  $e$  are distinct, and
3. there is a (possibly empty) set of distinct, actually occurring events  $\{d_1, d_2, \dots, d_n\}$  such that  $c$  influences  $d_1$ ,  $d_1$  influences  $d_2$ ,  $d_2$  influences  $d_3$ ,  $\dots$ , and  $d_n$  influences  $e$ .

By taking causation to be the ancestral of influence, Lewis has ensured that causation is always *transitive*. However, there seem to be intuitive counterexamples to the transitivity of causation, including:

- *Black and Red:* "Imagine a conflict between Black and Red. . . . Black makes a move that, if not countered, would have advanced his cause. Red responds with an effective countermove, which gives Red the victory. Black's move causes Red's countermove, Red's countermove causes Red's victory. But does Black's move cause Red's victory? Sometimes it seems not" (p. 194).
- *Birth and Death:* " 'Yeah, yeah, my birth is the cause of my death!' said the scoffer" (p. 196).

In response to Black and Red, Lewis tries to explain away the intuition that Black's move doesn't cause Red's victory. Two reasons that he cites why our intuitions might be confused about this case:

1. We might be confusing *what is generally conducive to what* with *what caused what*.
2. We might be confusing *whether-whether dependence* with *causation*.

In response to Birth and Death, Lewis argues that there are pragmatic reasons why we don't say that one's birth causes one's death: "His birth is indeed a cause of his death; but it is understandable that we seldom say so. The . . . dependence of his death on his birth is just too obvious to be worth mentioning" (p. 196).

## 2. Kagan on the Practice of Appealing to Intuitions about Cases

According to Kagan, we give central importance to our intuitions about cases (= "case-specific intuitions") when we argue for/against various moral claims:

- Case-specific intuitions are taken to provide reasons to accept/reject moral claims.
- We often try a moral theory "on for size" by considering its implications in a range of cases.
- Even those who officially disavow the legitimacy of appeals to case-specific intuitions usually make use of them in one way or another.
- When intuitions about general principles clash with intuitions about particular cases, we usually take the latter to have greater authority.

Although the extent to which we rely upon case-specific intuitions is widely recognized, Kagan doesn't think that we yet possess an adequate justification for our practice: "it is far from clear what, if anything, makes it legitimate for us to give these intuitions the kind of priority we typically give them" (p. 46).

Two sorts of vindications of intuitions about particular cases:

1. *A "deflationary" possibility*: the legitimacy of relying upon case-specific intuitions is simply an instance of the general legitimacy of relying on *any* belief (or disposition to believe) that we are confident in.
2. *A more ambitious possibility*: there is something special about case-specific intuitions which justifies our reliance on them (and moreover justifies the sort of priority we usually give them).

Most of us are drawn to this second view, says Kagan.

Kagan's article is specifically about *moral* intuitions. One thing to ask ourselves while assessing his claims:

- To what degree (if at all) does what Kagan say about moral intuitions apply *mutatis mutandis* to non-moral intuitions?

## 3. Kagan on the Analogy with Empirical Observation

Kagan thinks that the "closest we typically come" to justifying our reliance on moral intuitions is by exploiting a certain analogy:

*the analogy*: Case-specific intuitions play a role in the construction of *moral theories* analogous to that of *observations* in the construction of *scientific theories*.

"... it is precisely our moral views about examples, stories, and cases which constitute the data for moral theorizing" (Judith Jarvis Thomson, "Afterward" to *Rights, Restitution, and Risk*, p. 257).

In favor of the analogy with empirical observation:

- We give unique weight when constructing [scientific/moral] theories to accommodating our [empirical observations/case-specific moral intuitions].
- When [performing observations/eliciting intuitions], we can simply see—immediately and typically without further ado—that [the liquid in the test tube has turned red/a given act would be wrong].
- Any given [observation/intuition] can be rejected, but it is implausible to suggest that we should give no weight at all to our [observations/intuitions].
- Typically we give greater priority to preserving judgments based on [observation/intuition] than we do to maintaining allegiance to some [general empirical theory/general moral theory].

Potential points of disanalogy between moral intuition and empirical observation:

- i. *We have a fairly good idea of how empirical observations are produced, but not of how moral intuitions are produced.*

Empirical observations depend on the presence of well-functioning sense organs, but if there is a corresponding organ for moral intuitions—a “moral sense”—we know little about how it works. (Actually, this claim is less true now than it was in 2001.)

In fact, Kagan thinks the analogy would go through even if there is no such organ, and he thinks it is harmless to posit a moral sense as a placeholder for whatever it is that generates our intuitions.

- ii. *We have good reason to think that sense perception is reliable, but no reason (yet) to think moral intuition is reliable.*

Kagan thinks that the following justifies us in thinking our empirical observations to be reliable:

- a. “We find ourselves strongly inclined to believe these observations . . . and so in the absence of a good reason to reject them, it is reasonable to (continue to) accept them” (p. 49).
- b. “. . . we are able to incorporate these observations into an overall attractive theory of the empirical world, one which admittedly rejects some of the observations as erroneous, but which for the most part endorses the observations as correct” (ibid.).

Given the compelling nature of our case-specific moral intuitions, if we are able to construct an overall moral theory that accommodates most of those intuitions, then we could, in a similar way, justify our taking case-specific moral intuitions to be reliable.

Can we construct such a theory?

The theory must do more than organize our moral intuitions into systematic patterns: it must *explain* how the moral domain can have the features ascribed to it by our intuitions.

Kagan has argued elsewhere that *no* plausible moral theory can given this sort of theoretical underpinning to all—or even most—of our case-specific moral intuitions.

Kagan concludes that not only do we *lack* reason to believe that moral intuitions are *reliable*, but moreover we *have* reason to believe that they are *unreliable* (cf. p. 54).

- iii. *We have intuitions about both particular cases and general principles, whereas we only observe the former.*

One problem: intuition apparently has two kinds of object, whereas observation has only one.

A deeper problem: we give special priority to our intuitions about one of those kinds of object over our intuitions about the other kind, and this must be explained.

In fact, Kagan thinks that when intuition gives a verdict about a (so-called) particular case, really it is giving a verdict about a *type* of case. (This claim seems more difficult to defend for situations in which we have an intuition about an actual, concrete case that is before us.)

But if Kagan’s claim here is correct, then intuition is *always* applied to general principles, thus widening the gulf between intuition and observation.

- iv. *Unlike the case of observation, there is widespread intuitive disagreement about cases.*

Although he notes the difficulty in being sure whether this is so (for example: when we vote in polls, are we reporting our moral intuitions, or simply stating our tentative moral beliefs about a case?), Kagan thinks there is widespread disagreement among people’s case-specific moral intuitions. Furthermore, these differences seem to cluster into distinct patterns of intuitive response.

This disagreement does not seem analogous to *color blindness*: it is not clear how we can make good on the claim that one side is suffering from *moral blindness*, and moreover each side may be failing to respond to at least some features that the other side’s intuitions mark out as morally significant.

*Kagan's conclusion:* All of us must accept an *error theory* according to which at least many of our case-specific moral intuitions are mistaken.

However, this still leaves us with the task of providing such an error theory and explaining how, and in which ways, our reliance on intuition must be tempered.

Kagan himself is skeptical of *radical error theories* that dismiss most or all of our moral intuitions as suspect.

#### 4. Goldman & Pust's Attempt at Vindicating Intuition

A class M of contentful mental states is a *basic evidential source* if (and only if?) being in one of those states in favorable circumstances is *prima facie* evidence for the truth of its content.

- (RI) Contentful mental states of type M constitute a basic evidential source only if the M-states are *reliable indicators* of the truth of their contents (or the truth of closely related contents) when they occur in M-favorable circumstances.

Two other features standardly characterize basic evidential sources:

1. In M-favorable circumstances, the M-states *counterfactually depend* on the truth of their contents.
2. There is a distinctive *causal route* from the family of states of affairs that make the M-contents true or false to the family of M-states.

Goldman & Pust: “. . . if it is known or suspected that there is no relevant causal route or counterfactual dependence, there are grounds for doubting the existence of a reliable indicatorship relation” (p. 181).

(There is a worry that this claim constitutes an unwarranted bias against the possibility of there being basic evidential sources for knowledge of *mathematical truths*, *normative truths*, *necessary truths*, etc.)

Different ways of construing the *targets of philosophical analysis*:

- *universals* or *Platonic forms*;
- statements of *modal equivalence*;
- *natural kinds*;
- *concepts*<sub>1</sub>—concepts in the Fregean sense;
- *concepts*<sub>2</sub>—concepts in the personal psychological sense;
- *concepts*<sub>3</sub>—shared concepts<sub>2</sub>.

Views that embrace one of the first four possibilities Goldman & Pust call *extra-mentalistic approaches*, and views that embrace one of the last two they call *mentalist approaches*.

Goldman & Pust's basic complaint against extra-mentalistic approaches: on such approaches it is obscure how a causal relation or counterfactual dependence of the right sort could obtain, so we have good reason to doubt that the needed reliable indicatorship relation obtains.

Goldman & Pust endorse a mentalistic approach according to which the targets of analysis are concepts<sub>2</sub>.

This allows them to explain the reliable indicatorship relation: when I intuit *that case C counts as an instance of knowledge*, really I am intuiting *that C satisfies my concept<sub>2</sub> that I express with the predicate “knows.”*

However, this approach also seemingly undermines any interest we might have in the verdicts of intuition: who cares what my personal psychological concept of *knowledge*, *causation*, or *the good* is? Isn't the most crucial issue what personal psychological concept of these things I *should* have?

A general worry when considering various accounts of philosophical intuition:

How do we avoid the Scylla of embracing an account that renders the reliability of intuition a mystery, without falling into the Charybdis of accepting an account that can explain its reliability but only at the cost of rendering the deliverances of intuition philosophically uninteresting?