Causal Inference in Evidence-Based Policy. A tale of three monsters and how to defeat them

Alexander Gebharter Christian J. Feldbacher-Escamilla

Autumn 2020

Project Information

Talk(s):

- Gebharter, Alexander and Feldbacher-Escamilla, Christian J. (2020-10-02/2020-10-02). *Causal Inference in Evidence-Based Policy. A tale of three monsters and how to defeat them*. Public Lecture. Presentation (invited). Fellowship Lecture. IMTO University: Institute of Philosophy.

Project(s):

- DFG funded research unit *Inductive Metaphysics* (FOR 2495); subproject: *Creative Abductive Inference and its Role for Inductive Metaphysics in Comparison to Other Metaphysical Methods*.

Introduction

A few things we expect from a **good policy**:

- improve the overall situation
- no or little undesired side effects
- high efficacy
- cost/resource efficiency
- public support

Question: How can we predict the efficacy of a policy?

Example: Do face masks reduce spread of COVID-19?

Introduction

Two approaches to policy making:

Thinking vs. Evidence

Introduction

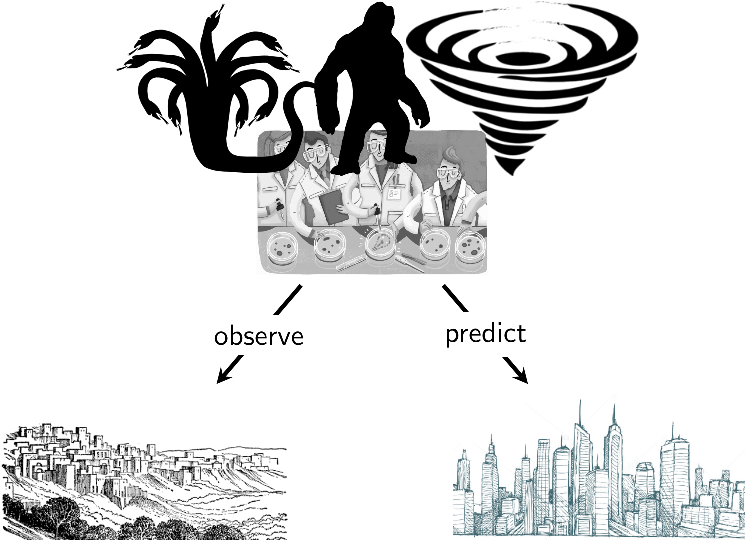

Contents

Evidence, Inference, and P-Kong

Evidence, Inference, and P-Kong

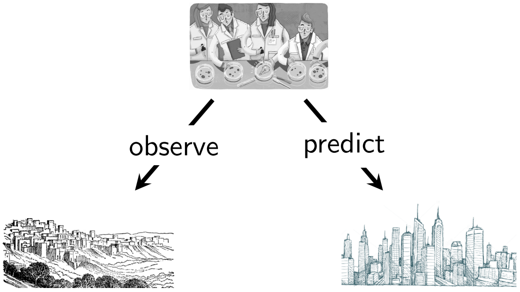

The **orthodox view**:

Evidence:

- Randomized control trials (RCTs)
- Meta studies

Inference (prediction):

- Induction

Evidence, Inference, and P-Kong

Randomized control trial (RCT):

- *Random* assignment of subjects into two groups:
    - Test group
    - Control group
- *Enforce* the policy (P) in the test group
- Compare the outcome (O) in the two groups

Upshot:

- If successful, P turns out to be an *effective means* to achieve O, because the RCT establishes P as a *cause* for O.
- We need causation to make this inference; mere correlation is not enough.

Explanation:

- Proper randomization guarantees that all the causal influences on subjects in the two groups of factors different from P are equal.
- Hence, any difference in O in the two groups must be due to P.

Evidence, Inference, and P-Kong

An Example: RCT to test a new "back to work" programme in a particular city (cf. Haynes et al. 2010, p. 9)

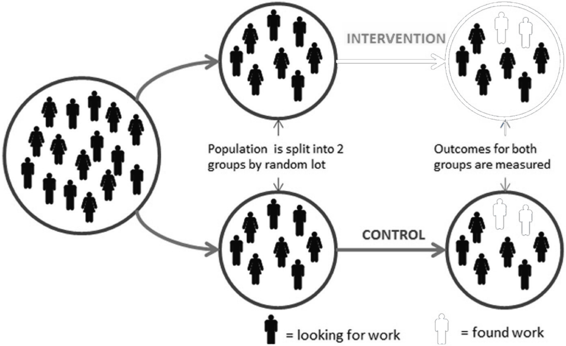

Randomization is key element: We can assume that the differences in the results are not due to differences between the groups, rather due to the intervention.

Evidence, Inference, and P-Kong

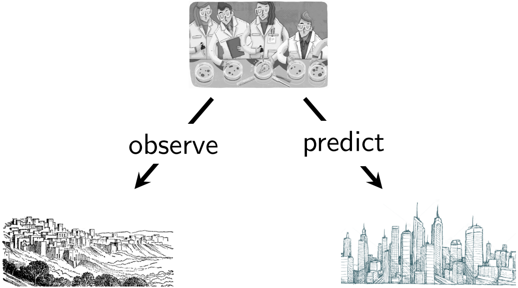

The **orthodox view**:

Evidence:

- Randomized control trials (RCTs)

Inference (prediction):

- Induction

Evidence, Inference, and P-Kong

Scientific Inference:

We can differentiate between three forms of scientific inference:

- **Deduction** . . . truth-preserving, explicative
- **Induction** . . . ampliative, but theoretically conservative
- **Abduction** . . . ampliative, but also theoretically innovative

Examples:

- We can deduce the Pythagorean theorem from elementary geometrical facts.
- We inductively infer that all swans are white based on our past observations of swans.
- We can abductively infer that it is gravitational influence of the Moon which causes the tides.

Evidence, Inference, and P-Kong

Induction:

- An inference method that generalizes n observations that policy P worked to P also working for case $n+1$.

General shema:

$$\begin{array}{l}
\text{Policy } P \text{ worked in city } 1. \\
\vdots \qquad\qquad\qquad\qquad\quad \vdots \\
\text{Policy } P \text{ worked in city } n. \\
\hline
\text{Policy } P \text{ will work in city } n+1.
\end{array}$$

Particularly **Karl Popper** stressed: Induction is prone to error.

Evidence, Inference, and P-Kong

An Example:

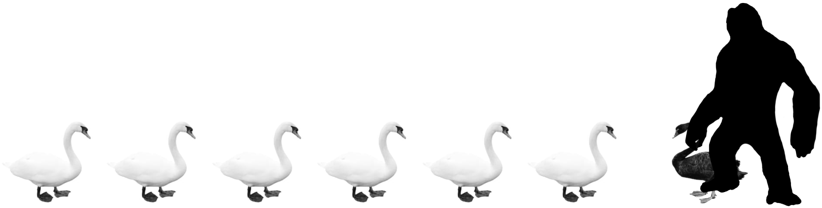

So, given P worked in case $1, \ldots, n$, does NOT provide any guarantee that P works also for all cases/for case $n + 1$.

Popper: We can only infer (given this data by deduction):
P works NOT in all cases.

Evidence, Inference, and P-Kong

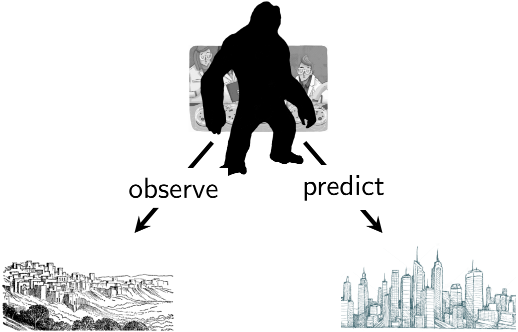

Cartwright & Hardie (C&H): The orthodox view is threatened by Popper Kong (P-Kong).

- The RCT only shows that *P* worked in city 1 (with a specific causal profile).
- *P* might not work in city 2 (with a different causal profile)

Cartwright & Hardie on Defeating P-Kong

Cartwright & Hardie on Defeating P-Kong

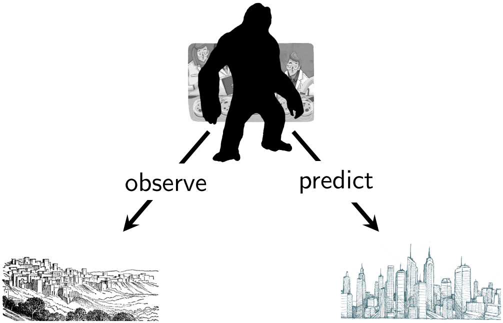

The **orthodox view**:

Evidence:

- Randomized control trials (RCTs)

Inference (prediction):

- Induction ⇐ Culprit according to C&H

Cartwright & Hardie on Defeating P-Kong

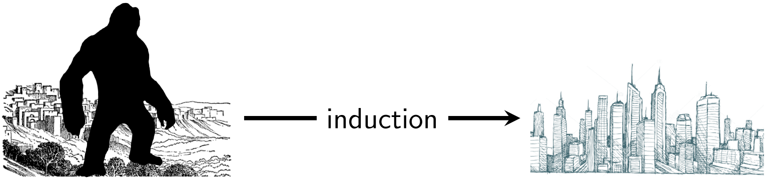

C&H: Policies are like a special ingredient in a **cake**; it only works if the other ingredients (**support factors** F) are right.

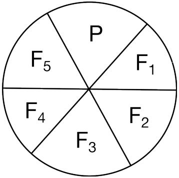

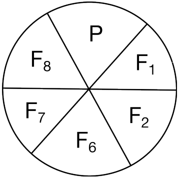

Cartwright & Hardie on Defeating P-Kong

C&H: Replace induction by **deduction** (an argument where the truth of the conclusion is necessitated by the premises):

General shema:

P worked in city 1.
The same *support factors* for *P* in city 1 are also present in city 2.
P plays the same *causal role* in city 1 as it played in city 2.

P will work in city 2.

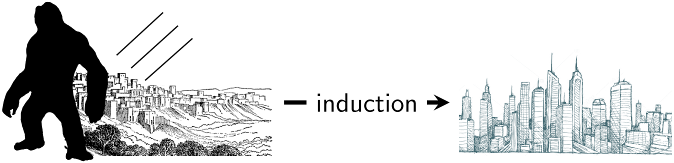

Upshot: We also have to think about support factors and causal roles in city 2.

New Monsters: Skylla & Charybdis

New Monsters: Skylla & Charybdis

Now there are **2 possibilities**:

- We fully know P's support factors and causal profile in city 2.
- We do not fully know P's support factors and causal profile in city 2.

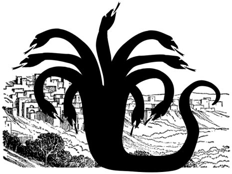

— deduction →

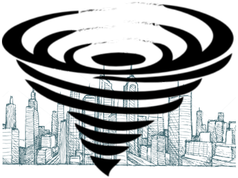

New Monsters: Skylla & Charybdis

Case 1: We fully know P's support factors and causal profile in city 2.

This means:

- We know city 2's causal cake, and
- we know whether all the relevant ingredients are present in city 2.

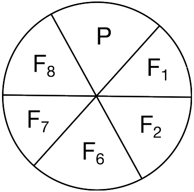

Skylla: City 1's causal cake (evidence) becomes irrelevant for inferring P's efficacy in city 2. $\Rightarrow$ Undermines whole idea behind evidence-based policy!

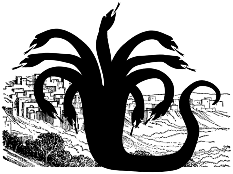

— deduction →

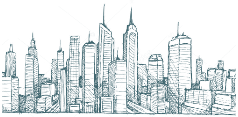

New Monsters: Skylla & Charybdis

Case 2: We do not fully know P's support factors and causal profile in city 2.

This means:

- We do not know city 2's causal cake, or
- we do not know whether all the relevant ingredients are present in city 2.

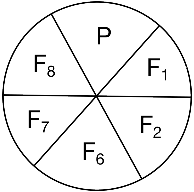

Charybdis: P's efficacy in city 2 can only be inferred on the basis of city 1's causal cake (evidence) inductively. $\Rightarrow$ Opens the gates for P-Kong!

— induction →

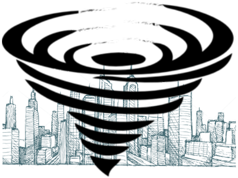

New Monsters: Skylla & Charybdis

Summarizing: C&H's deductive account attracts

- Skylla if we possess all the information to infer P's efficacy in city 2, or
- Charybdis (and, thus, P-Kong) otherwise.

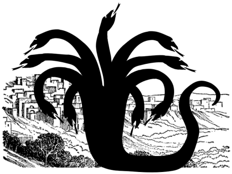

— deduction →

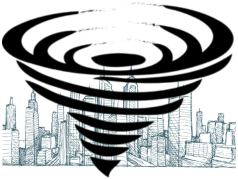

In any case, city 2 will be a mess!

Causal Inference to the Rescue

Causal Inference to the Rescue

How can we chase away all three monsters?

Proposal:

- Do not focus so much on how an efficacious policy in city 1 can be copied to city 2.
- Rather, try to learn the overarching causal structure responsible for the success/failure of P in different cities $1.1 - 1.n$.

So the inference pattern we want is not:

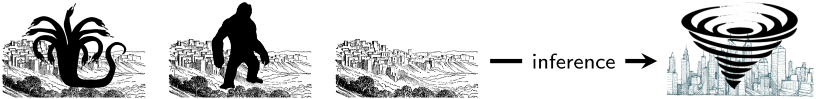

Causal Inference to the Rescue

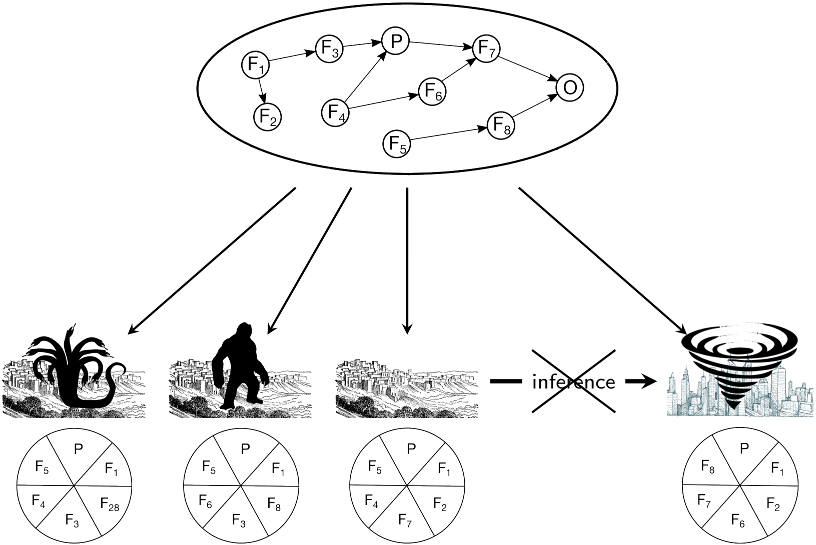

Causal Inference to the Rescue

We proceed in **4 steps**:

- Infer the overarching causal structure S best explaining *all* the evidence in cities 1.1-1.n.
- Test and improve S.
- Observe as many of S's factors F_i in city 2 as possible.
- Use these observed factors together with S to predict whether and to what extent P will be efficacious in city 2.

Note:

- Steps 1 (and 2) involve abductive inference and require creativity.
- S allows for novel predictions and can be tested independently.
- The more factors S involves and the better it is confirmed, the more reliably it is able to predict P's efficacy in city 2.
- Thus: By expanding and confirming S, we increase the likelihood of P's efficacy in city 2.

Causal Inference to the Rescue

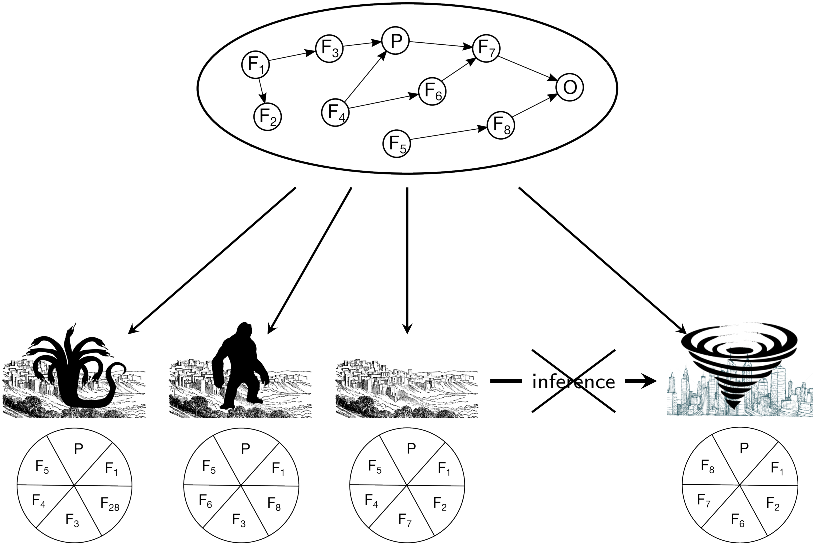

Causal Inference to the Rescue

Our diagnosis fits into a more general pattern as outlined by the champion of the probabilistic approach to artificial intelligence.

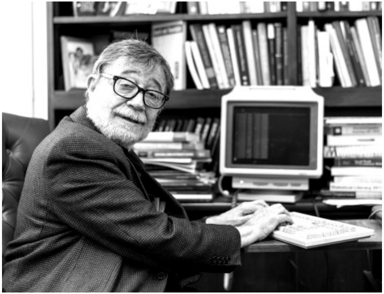

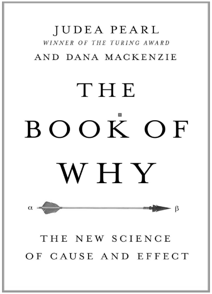

The "Causal Revolution" in AI: We no longer aim at describing WHAT is the case, but also: WHY it is the case.

> *"[R]eturning the Causal Revolution to its womb in artificial intelligence, I aim to describe to you how robots can be constructed that learn to communicate in our mother tongue—the language of cause and effect." (Pearl 2018)*

Causal Inference to the Rescue

The Causal Revolution in AI:

Orthodox Statistical Analysis		**RCTs**		**Causal Inference**
correlation-centred observational studies	⇒	study of "provisional causality"	⇒	full-blown study of causation

Our investigation of C&H: causal inference is also key for AI-based or AI-assisted policy making.

Summary

Our investigation shows:

- Simple inductive and deductive reasoning does not suffice for good policy.
- We need more powerful tools from AI (esp. **causal modeling**) in order to:
  - Form causal hypotheses on the basis of observational & experimental data.
  - Generate predictions about what would happen if factors were distributed such and such that form the basis for testing causal hypothesis.
  - Can generate predictions about what would happen under hypothetically possible policy interventions in different causal contexts.

Selected References I

Cartwright, Nancy and Hardie, Jeremy (2012). *Evidence-Based Policy: A Practical Guide to Doing It Better*. Oxford: Oxford University Press.

Haynes, Laura, Service, Owain, Goldacre, Ben, and Torgerson, David (2012). "Test, Learn, Adapt: Developing Public Policy with Randomised Controlled Trials". In: *Cabinet Office - Behavioural Insights Team*. DOI: 10.2139/ssrn.2131581.

Pearl, Judea and Mackenzie, Dana (2018). *The Book of Why. The new science of cause and effect*. New York: Basic Books.