Bayesian data analysis & cognitive modeling

Session 12: Bayesian ideas in philosophy of science

Michael Franke

Philosophy of science

goal

a theory of science

descriptive

how science *actually is*

normative

how science *should be*

Some provocative questions

1. What is (or should be) the goal of scientific inquiry?
2. How do (or should) scientists try to achieve this goal?
3. What role does statistical inference play in science?
4. Which one promises to be more naturally conducive to the goal of science: Bayesian inference or NHST?

overview

Philosophy of science

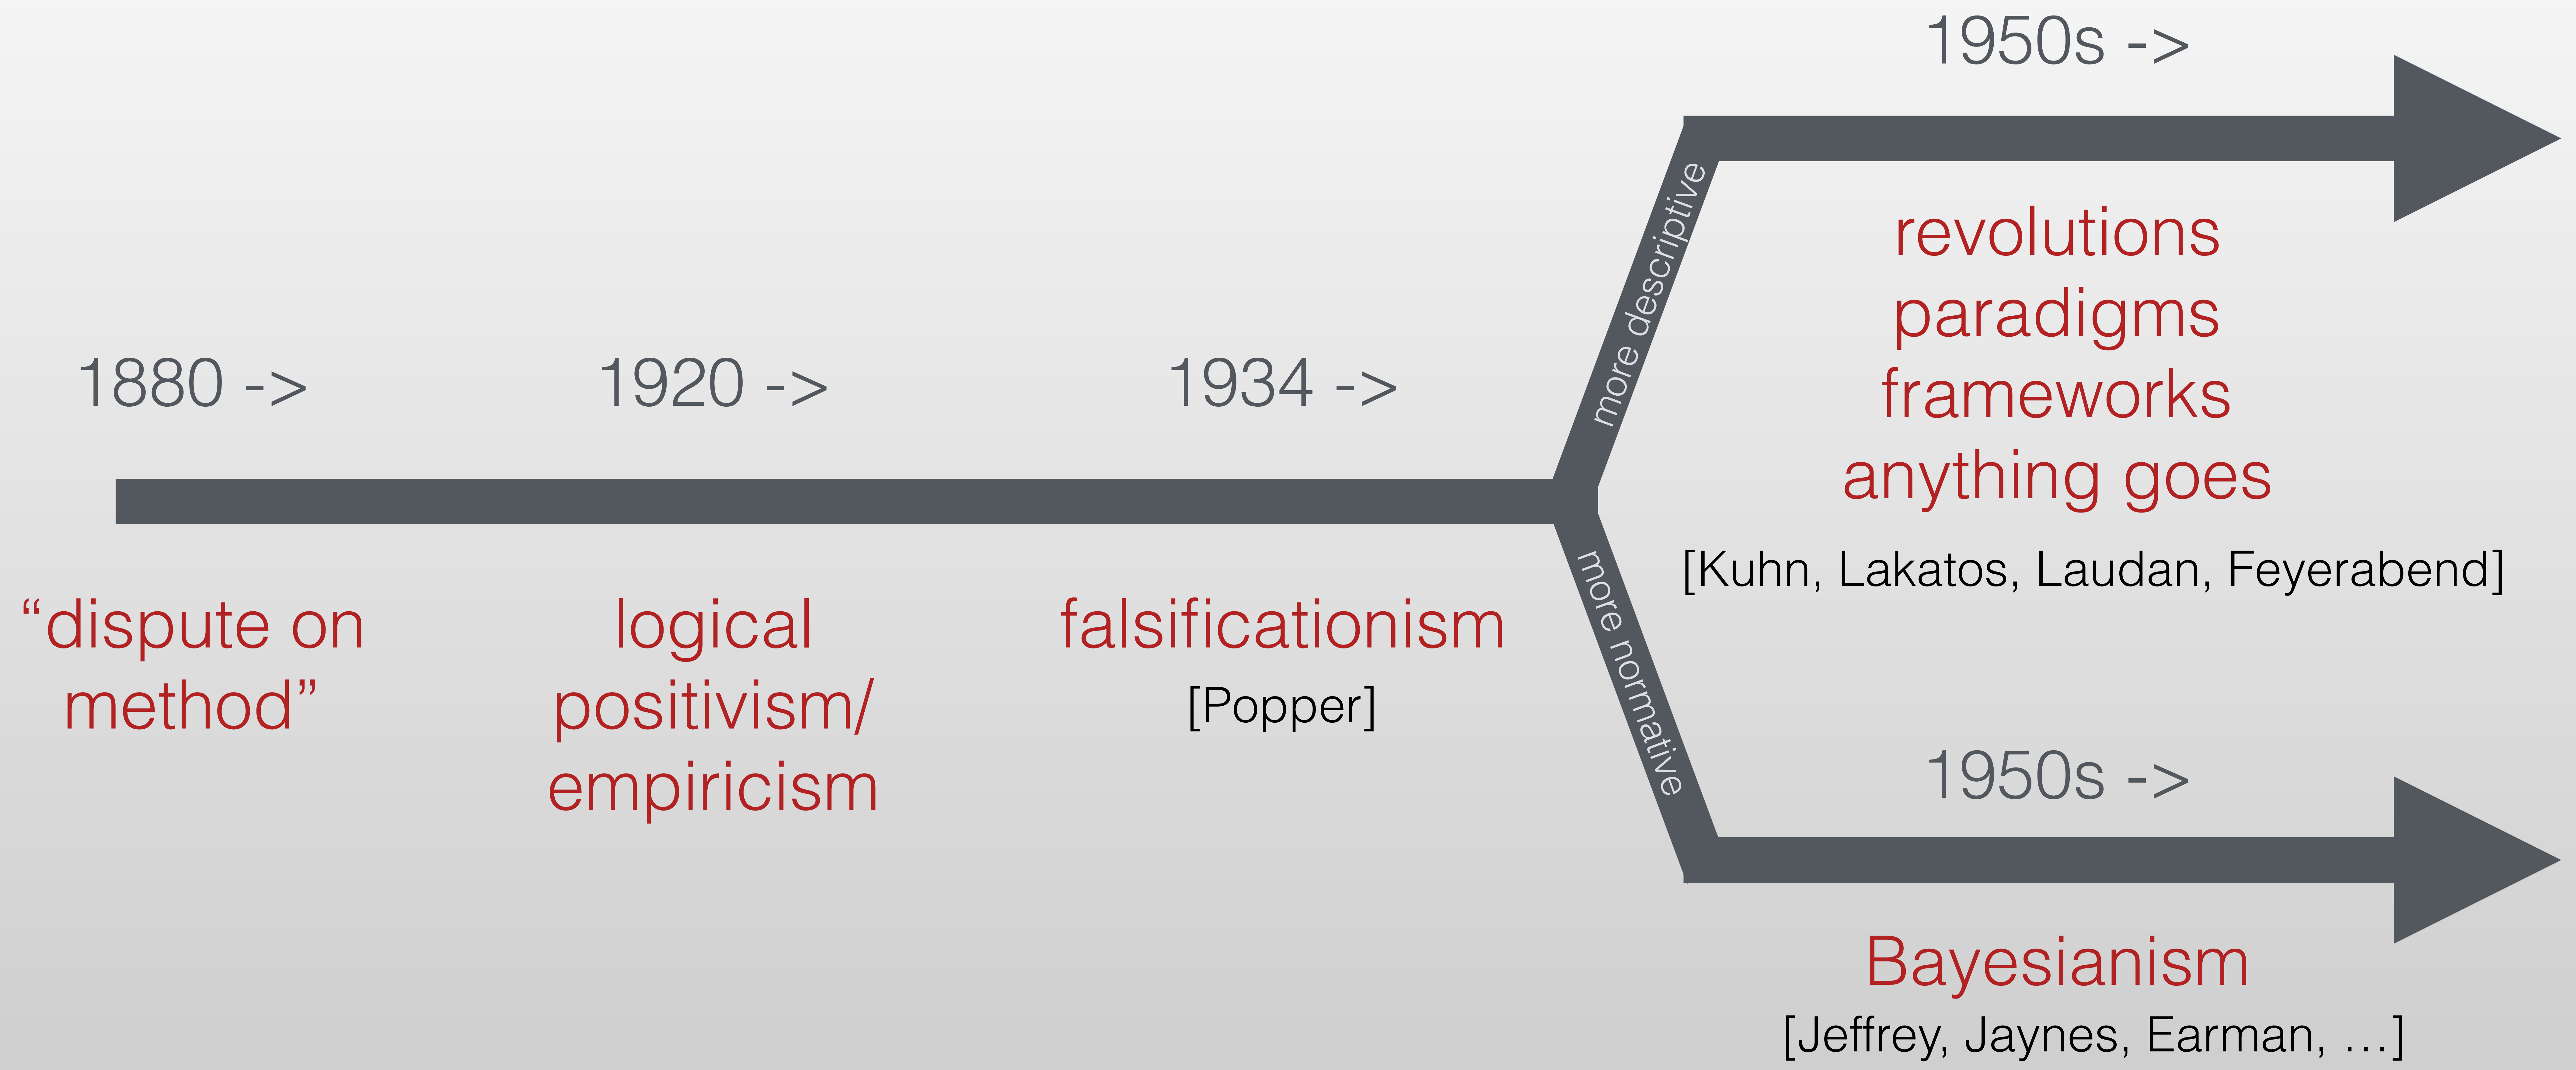

Crucial notions

falsification

explanation

George Washington Carver, botanist

confirmation

prediction

evidence

Popper:

demarcation & falsifiability

Sir Karl Raimund Popper

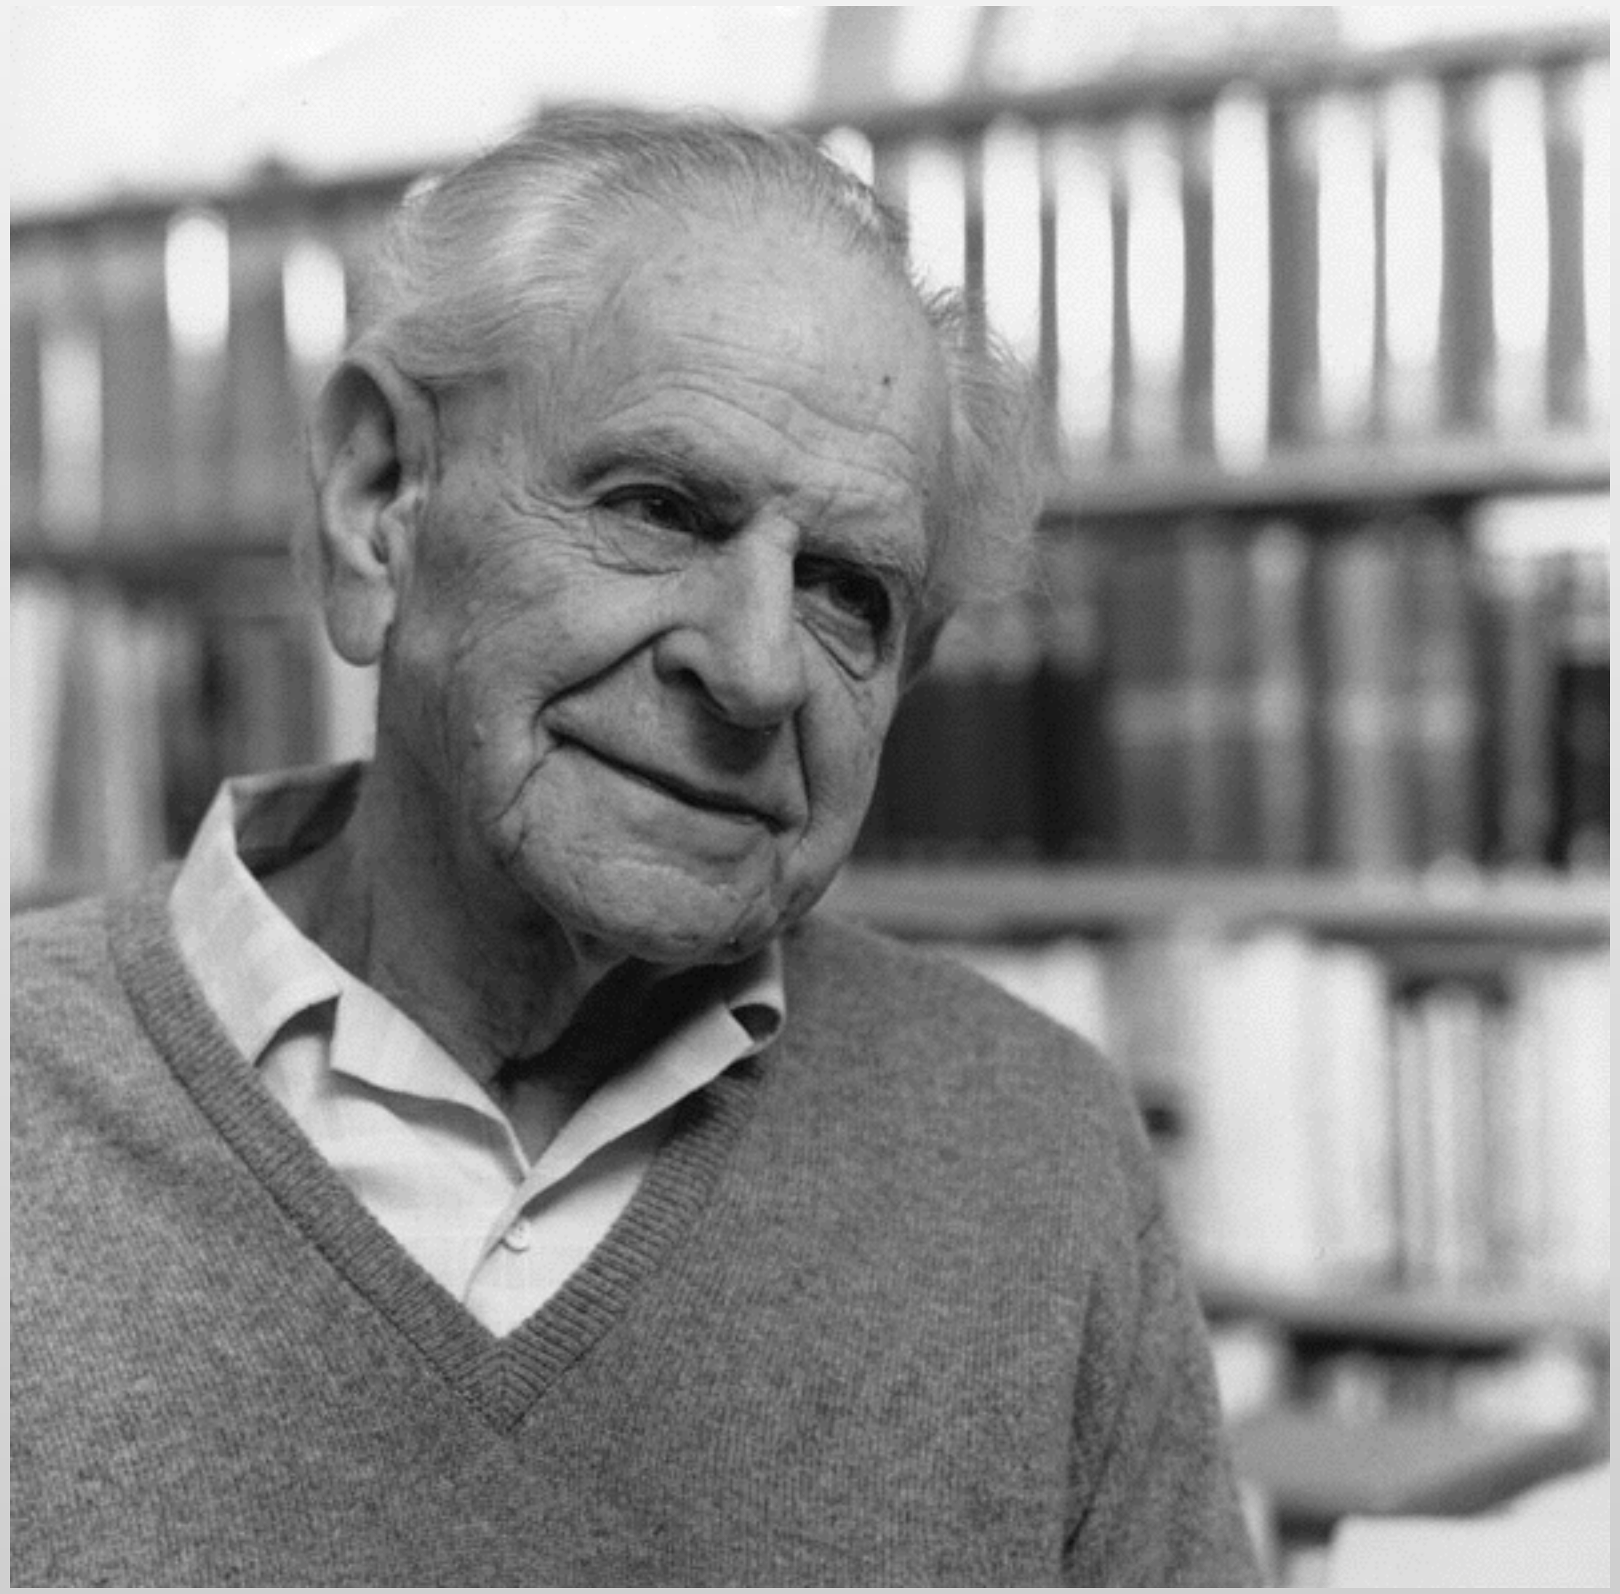

life & thought

born 28 July 1902 in Vienna
critical exchange with Vienna circle
emigrated to New Zealand during WW2
reader & professor in London (LSE)
influential work: “Logik der Forschung” (1934)
died 17 September 1994 in Kenley (London)

Main themes

goal: demarcation

distinguish science (Einstein) from pseudo-science (Marx, Freud)

solution: falsifiability

hypothesis *h* is scientific iff it has the potential to be falsified by some possible observation

falsification

hypothesis *h* is falsified if it logically entails *e* and we observe not-*e*

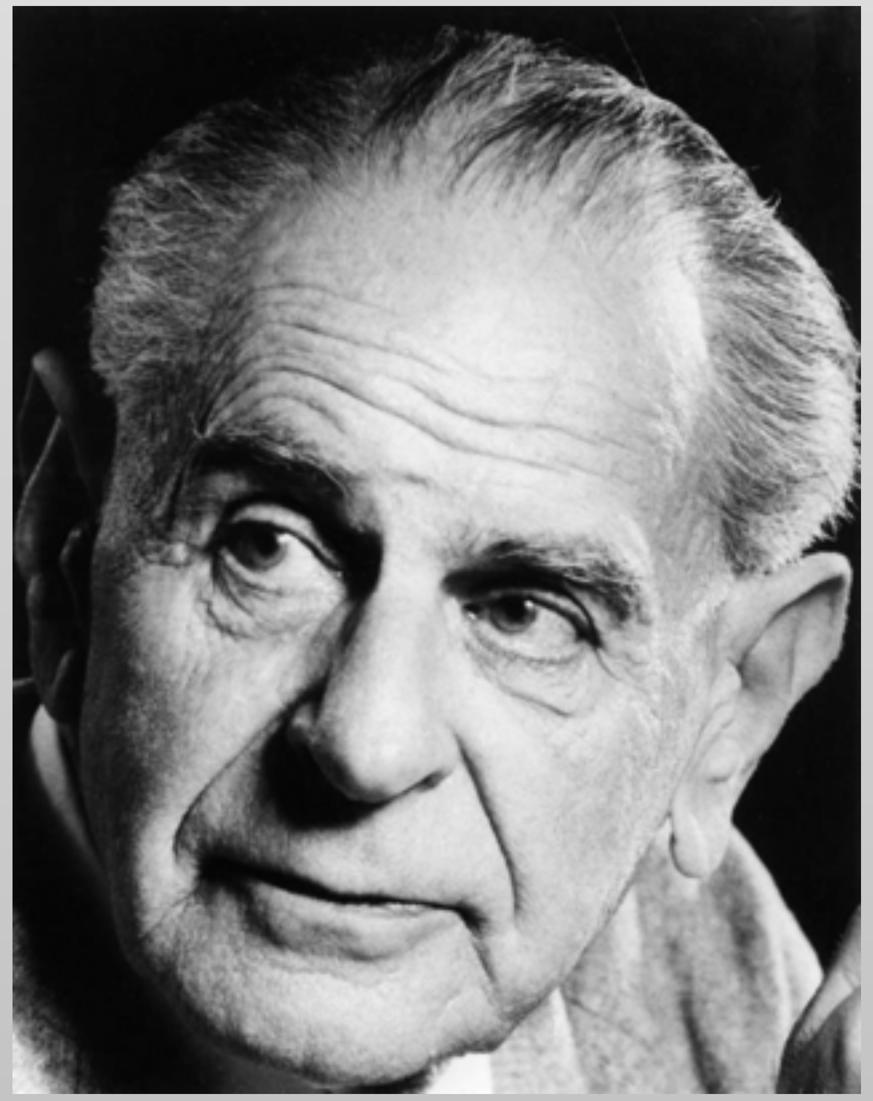

anti-confirmationism, fallibilism & tentativism

hypothesis *h* can never be confirmed by empirical evidence
hypothesis *h* is never 100% certain
maintain hypothesis *h* until refuted by evidence

Theory change

conjecture

refutation attempt

actual Popperian

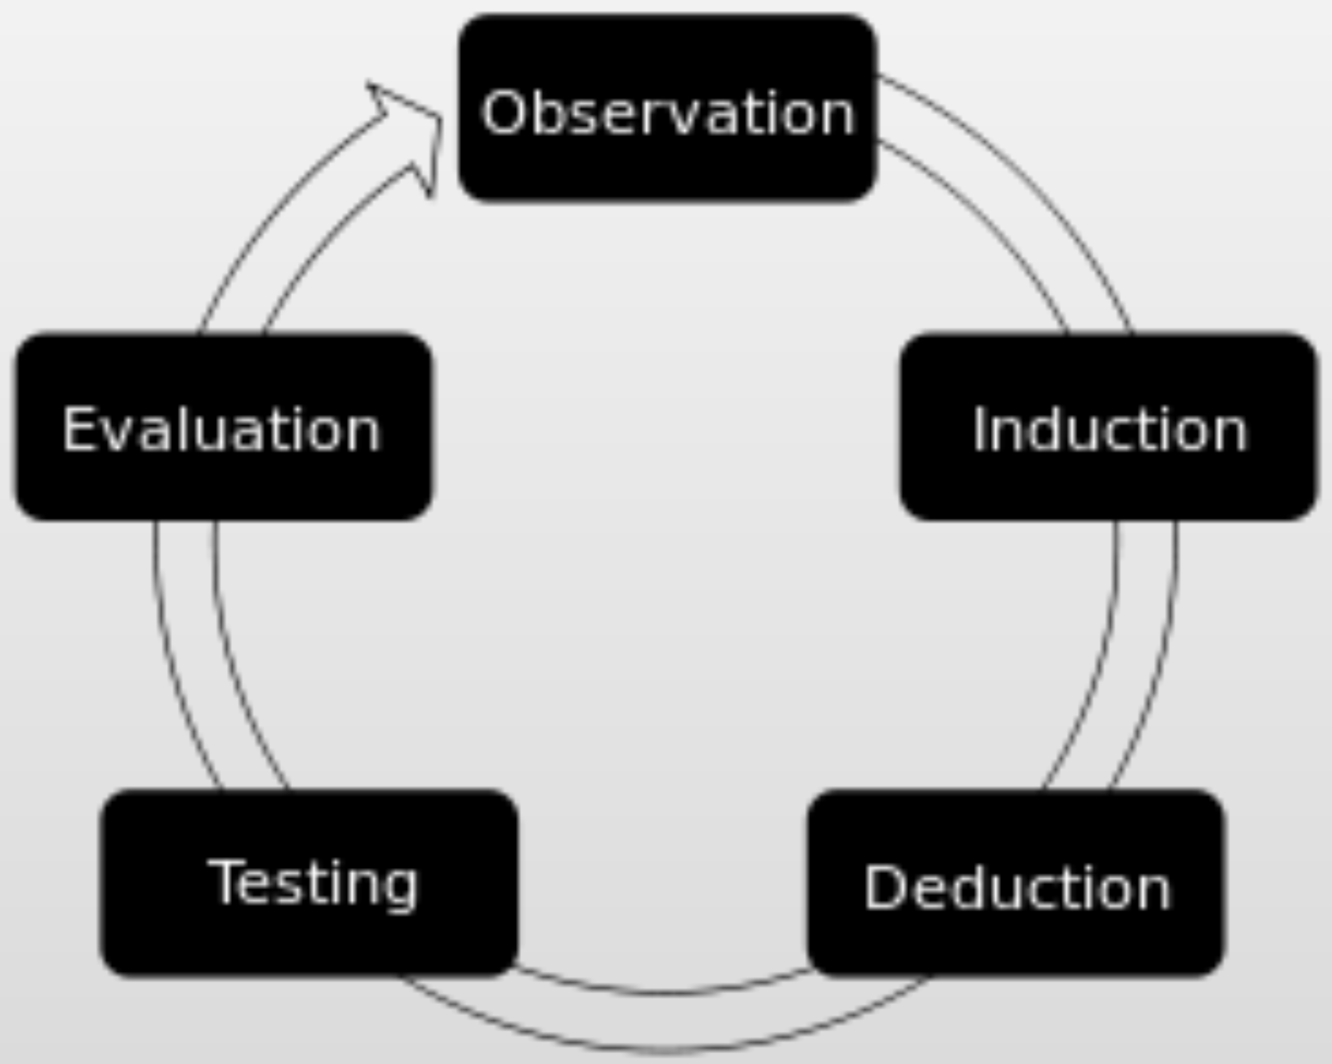

modern Popperian

how to form new conjectures? — be bold!

new hypotheses should make sharp predictions & increase breadth of applicability

Problems with falsifiability

holism of testing

when does *e* falsify *h* beyond any doubt?

- Quine-Duhem: can only test conjunction of “core theory” + “auxiliary assumptions”
- Popper: good scientist blames “core theory”

probabilistic predictions

what if *h* only makes certain observations unlikely, not logically impossible?

- Popper: not a scientific theory

Problems with anti-conformationism

practical decision making

why use currently adopted *h* and not arbitrary (untested *h'*) for practical applications?

- Popper: notion of "corroboration" (not "confirmation")
  - *h* is more corroborated the more refutation attempts it survived
- common sense: the more predictions of *h* come out correct, the likelier *h* appears

“only a theory”

video

Against anti-conformationism

considering positive evidence in favor of a theory is:

- natural
- essential for practical decision making
- important for deflecting the anti-scientific ”just a theory” farce

In defense of a weak falsificationism

Attitudinal Popperianism

demarcation of scientific attitude from unscientific attitude

- it matters less whether *h* is scientific or not (as a formal construct)
- it matters more whether we approach *h* in a “scientific manner”
  - formulate *h* as precisely as possible so that implications are clear
  - try to check implications empirically
  - never mistake *h* for fact (fallibilism: “only a theory”)

do not reject ideas for what they are,
reject attitudes towards critical assessment of ideas

Null-hypothesis significance testing

researchers celebrating p=0.048

Popper vs NHST

Popper's falsificationism

look for observations that would likely falsify current hypothesis/theory H_1

NHST in usual practice

according to H_1 we predict an effect (e.g., difference or means…)
H_0 assumes absence of effect
significant *p*-value ==> reject H_0
treated as support for H_1

Bayesianism

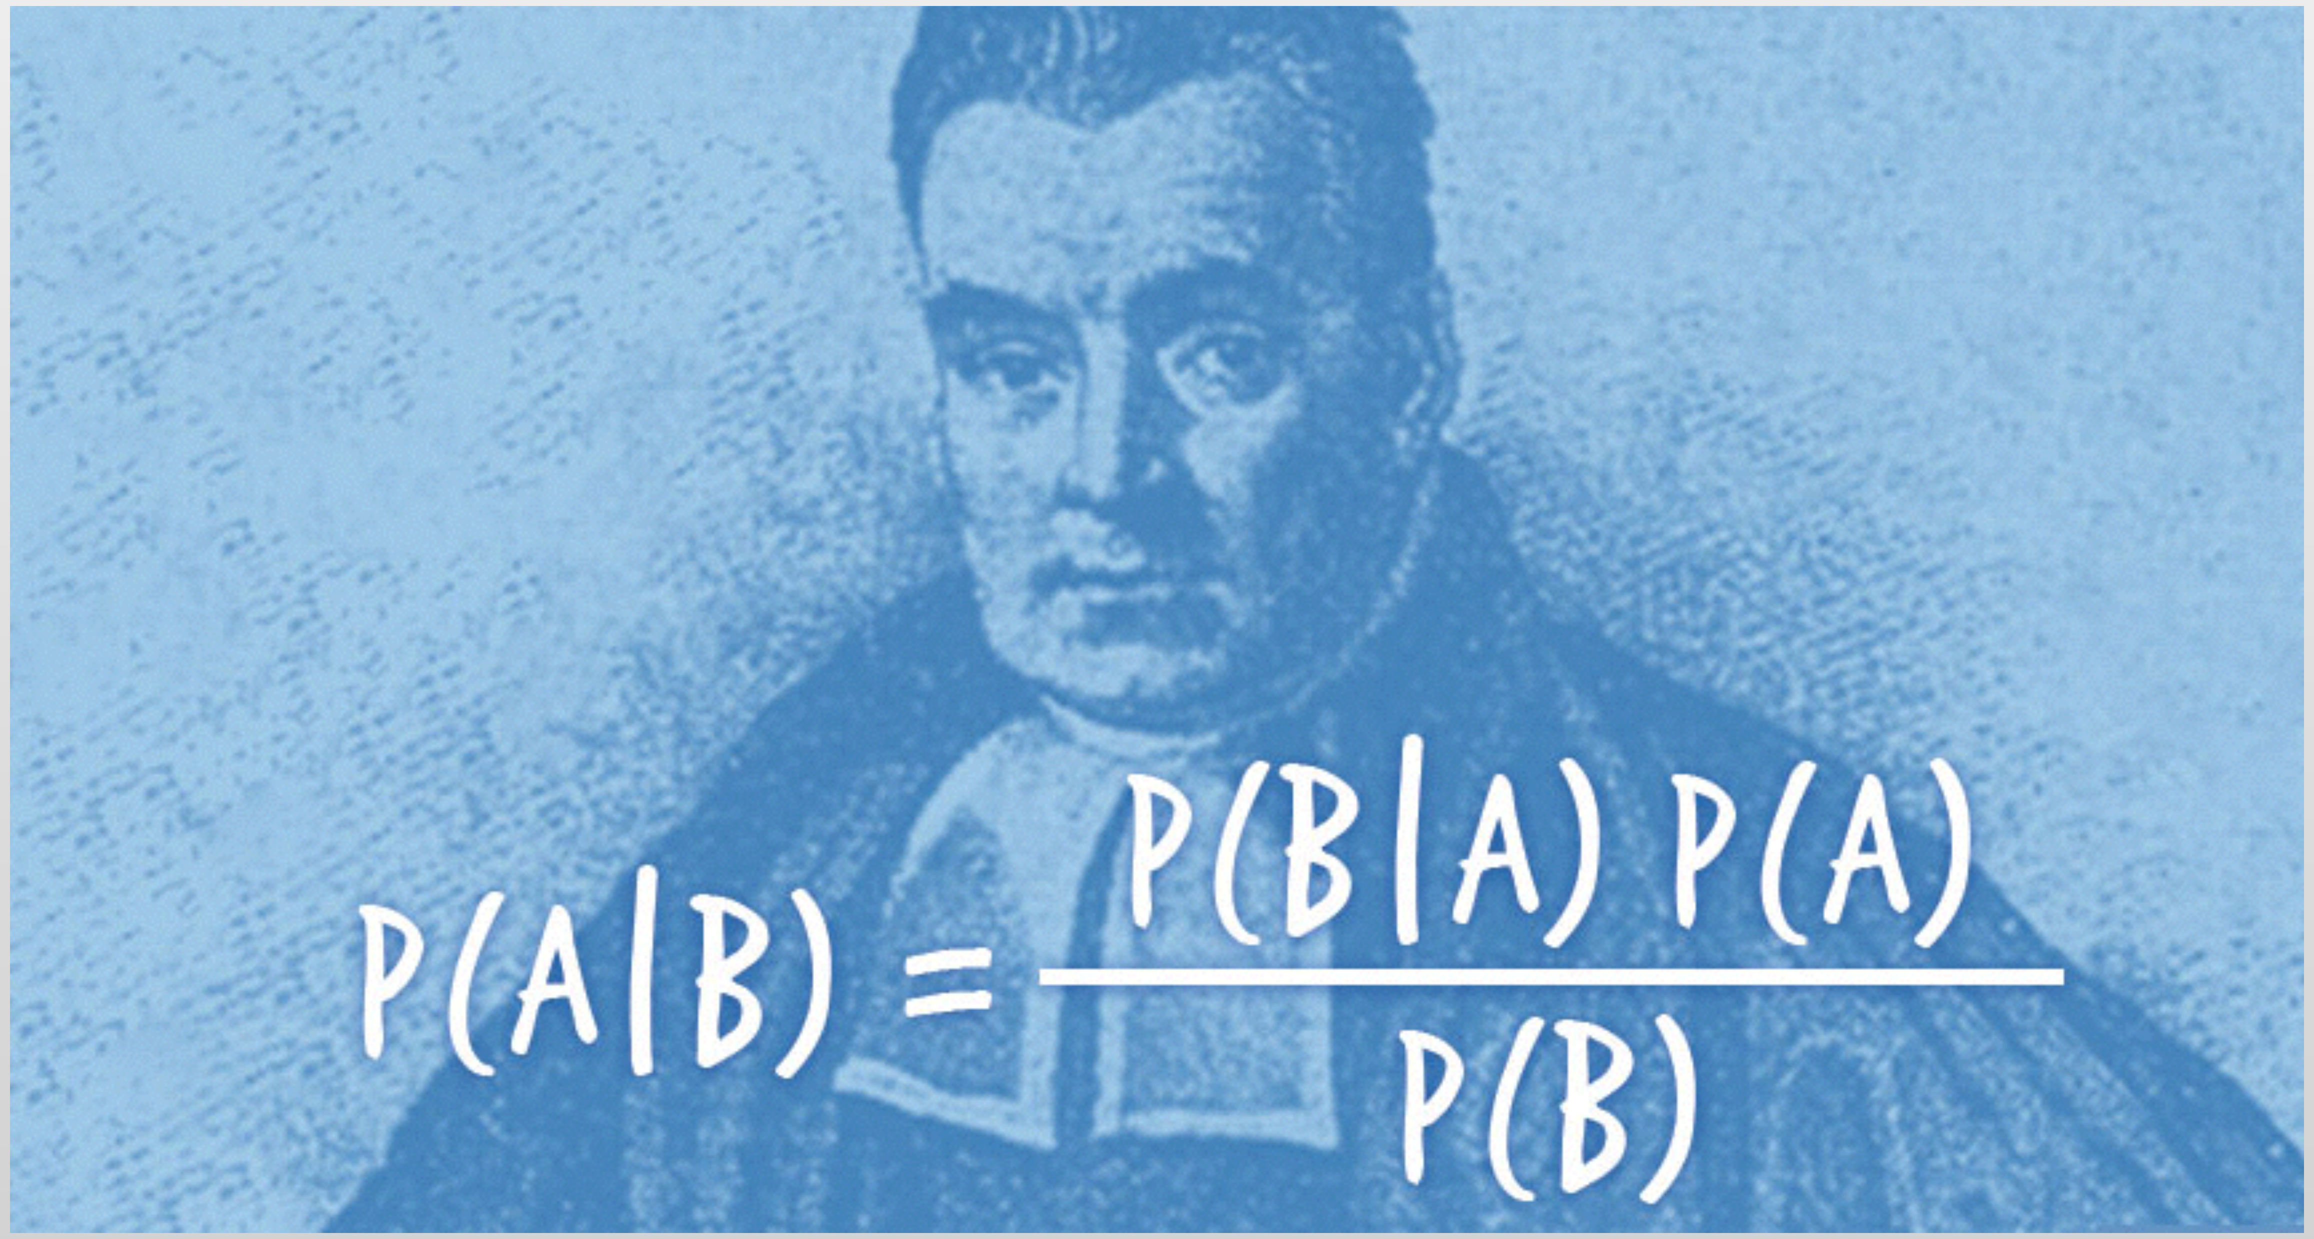

First-shot formalizations from a Bayesian point of view

confirmation & evidence

observation *e* <u>confirms</u> hypothesis *h* if *e* is/provides positive <u>evidence</u> for *h*

[i.o.w., confirmation is absolute where evidence is quantitative]

e is/provides positive <u>evidence</u> for *h* if *h* is made more likely by *e*

$$P(h \mid e) > P(h)$$

explanation

hypothesis *h* <u>explains</u> observation *e* if *h* makes *e* less surprising

$$P(e \mid h) > P(e)$$

prediction

hypothesis *h* <u>predicts</u> observation *e* if *e* is expectable under *h* but not otherwise

$$P(e \mid h) > P(e \mid \bar{h})$$

Bayesian evidence

evidence

e is/provides positive evidence for h if h is made more likely by e $P(h \mid e) > P(h)$

by Bayes rule & expansion

$$P(h \mid e) = \frac{P(e \mid h)P(h)}{P(e)} = \frac{P(e \mid h)P(h)}{P(e \mid h)P(h) + P(e \mid \overline{h})P(\overline{h})}$$

frequently raised problem

need to know likelihoods $P(e \mid h)$ and $P(e \mid \text{not-}h)$, as well as priors $P(h)$ and $P(\text{not-}h)$

Bayesian evidence

not so

$$P(h) < P(h \mid e)$$

$$P(h) < \frac{P(e \mid h)P(h)}{P(e)}$$

$$P(h) < \frac{P(e \mid h)P(h)}{P(e \mid h)P(h) + P(e \mid \overline{h})P(\overline{h})}$$

$$P(e \mid h) > P(e \mid h)P(h) + P(e \mid \overline{h})P(\overline{h})$$

$$P(e \mid h) > P(e \mid h)(1 - P(\overline{h})) + P(e \mid \overline{h})P(\overline{h})$$

$$P(e \mid h) > P(e \mid \overline{h}) \quad [\text{if } P(\overline{h}) \neq 0]$$

upshot

observation *e* is evidence for hypothesis *h* if *e* is more likely under *h* than under *not-h*
=> only likelihoods required

relation to Bayes factors

strength of evidence is a function of how much bigger P(e|h) is than P(e|¬h)

a Bayesian notion of "explanation"

same story

$$P(e \mid h) > P(e)$$

$$P(e \mid h) > P(e \mid h)P(h) + P(e \mid \overline{h})P(\overline{h})$$

$$P(e \mid h) > P(e \mid h)(1 - P(\overline{h})) + P(e \mid \overline{h})P(\overline{h})$$

$$P(e \mid h) > P(e \mid \overline{h}) \quad [\text{if } P(\overline{h}) \neq 0]$$

upshot

h explains *e* iff *e* is evidence for *h*
=> only likelihoods required

Same same, but different (perspective)

confirmation & evidence

observation e confirms hypothesis h if e is/provides positive evidence for h
[i.o.w., confirmation is absolute where evidence is quantitative]
e is/provides positive evidence for h if h is made more likely by e

$$P(h \mid e) > P(h)$$

explanation

hypothesis h explains observation e if h makes e less surprising

$$P(e \mid h) > P(e)$$

prediction

hypothesis h predicts observation e if e is expectable under h but not otherwise

$$P(e \mid h) > P(e \mid \bar{h})$$

Pros and cons of Bayesianism

pro

intuitive quantitative formalization of evidence (e.g., Bayes factor)

no problems with theories that “just” make probabilistic predictions

seamless integration of uncertainty about auxiliary assumptions (think: Quine-Duhem problem)

does not require priors $P(h)$

con

requires likelihood functions $P(e|h)$

requires complete space of all relevant theories for $P(e|\neg h)$ or is necessarily relative to subset of graspable theories