Induction and confirmation

Christian Wüthrich

http://philosophy.ucsd.edu/faculty/wuthrich/

15 Introduction to Philosophy: Theory of Knowledge
Spring 2010

Sir Francis Bacon (1561-1626)
Induction in scientific practice

- English philosopher, scientist, statesman, lawyer, author
- systematic establishment and popularization of inductive methodology
- argued that reasoning from fact to axiom to law must be inductive, rather than deductive (as was the case in the Aristotelian tradition)

$\Rightarrow$ induction enters canon of scientific enquiry as the basis of scientific method

Types of non-deductive inferences

1. induction narrowly construed or enumerative induction: Swan 1 observed at time t_1 was white, swan 2 observed at time t_2 was white... $\Rightarrow$ All swans are white
2. projection: Swan 1 observed at time t_1 was white,... swan $n-1$ observed at time t_{n-1} was white $\Rightarrow$ Swan n (the next one to be observed) will be white
3. abduction, inference to the best explanation, or explanatory inference: data $\Rightarrow$ hypothesis about a structure or process that would "explain" the data

debate about which non-deductive inference is most basic: Hans Reichenbach (induction) vs. Gilbert Harman (abduction)

Nota bene: I will use the term "induction" to refer to the first two types of inference above.

The Mother of All Problems...

Problem (The problem of induction)

The reliability of an inductive inference from past experience to prediction concerning future must be underwritten by a principle of induction. But such a principle cannot claim logical necessity; nor can it be based on the past success of induction on pain of circularity. So how could the use of this principle be justified?

David Hume's problem of induction in brief

Hume, *Enquiry*, Section V

- relations of ideas vs. matters of fact
- relations of ideas: can be known independently of observation, abstract realm of logic and mathematics, all analytic *a priori* beliefs
- matters of fact: everything that is not a relation of ideas, concerns material existence, synthetic knowledge
- matters of fact can be observed (e.g. "there is a desk here"), or unobserved (e.g. "the sun will rise tomorrow")
- In order to know any matter of fact beyond what is directly given by our sensory experience, inductive reasoning must be employed.

- Inductive inferences depend on a "principle of the uniformity of nature" (or principle of induction): past acts as reliable guide to the future.
- Hume argues that such a principle cannot be justified; rational justification, were we to have it, could come in two different forms:
    1. demonstrative, *a priori* reasoning; but future does not depend logically on past bc it is conceivable that future does not resemble past; cannot ground induction in *a priori* reasoning
    2. inductive reasoning: our past success in using inductive inference warrants inductive inferences into the future; circular!

$\Rightarrow$ Conclusion: inductive practices have no rational foundation.

A bit more detail...

Huemer 293f

1. "Our beliefs can be divided into three categories: (a) Beliefs about relations of ideas... Beliefs about observed matters of fact... (c) Beliefs about unobserved matters of fact...
2. "All unobserved matter-of-fact beliefs depend upon inductive inference for their justification...
3. "All inductive inferences presuppose some such premise as 'The course of nature is uniform' or 'Unobserved things will resemble observed things.' Call this the 'Uniformity Principle.'
4. "The Uniformity Principle is not a relation of ideas proposition, since it is not analytically true. It is logically consistent to hypothesize that the course of nature may not be uniform.
5. "The Uniformity Principle is not an observed matter of fact, since it makes a claim about unobserved objects.

6. "The Uniformity Principle is an unobserved matter of fact belief. (From 1, 4, 5)
7. "The Uniformity Principle depends for its justification on induction. (From 2, 6)
8. "But the Uniformity Principle cannot be justified by induction, since all inductive inferences presuppose the Uniformity Principle (premise 3), and circular reasoning is not acceptable.
9. "The Uniformity Principle cannot be justified. (From 7, 8)
10. "No inductive inference can be justified. (From 3, 9)"

C D Broad (1926):
induction is the glory of science and the scandal of philosophy...

Hume's solution

Hume, *Enquiry*, Section V

- justification of induction vs. description of practice of inductive inference
- solution of problem of explaining induction: nature
- inevitability of induction must be accepted
- Hume: if you insist on sound deductive justifications for everything, then you will starve to death: you wouldn't assume that bread nourishes you based on past experience
- almost evolutionary account for our inductive tendencies $\Rightarrow$ naturalism
- So what's the "Principle of Induction"?

"This principle is Custom or Habit. For wherever the repetition of any particular act or operation produces a propensity to renew the same act or operation, without being impelled by any reasoning or process of the understanding, we always say, that this propensity is the effect of Custom. By employing that word, we pretend not to have given the ultimate reason of such a propensity. We only point out a principle of human nature, which is universally acknowledged, and which is well known by its effects...

*"All inferences from experience... are effects of custom, not of reasoning... Custom, then, is the great guide of human life. It is that principle alone which renders our experience useful to us, and makes us expect, for the future, a similar train of events with those which have appeared in the past. Without the influence of custom, we should be entirely ignorant of every matter of fact beyond what is immediately present to the memory and senses. We should never know how to adjust means to ends, or to employ our natural powers in the production of any effect. There would be an end at once of all action, as well as of the chief part of speculation." (*Enquiry*, Section V, Part I, 307f)*

Sir Karl Popper (1902-1994)

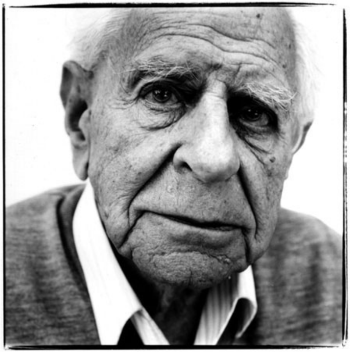

- born in Vienna, educated at U of Vienna
- 1928 PhD, 1930-1936 secondary school teacher
- 1934 *Logik der Forschung* (tr. 1959)
- 1937 emigration to NZ, lecturer at Canterbury U College of NZ
- 1946 emigrated to UK, position at LSE
- 1963 *Conjectures and Refutations*
- popular in science; "Popperazzi"

Popper's theory of science in a nutshell

- solution of prob of induction: forget about induction altogether and replace it with deductive method of testing
- "deductivism" (as opposed to inductivism)
- scientific progress results from the continued cycles of "conjectures" and "refutations"
- crucial: tentative attitude toward scientific thys
- science is search for truth, but we can never know whether we attained it!
- can never be completely sure that a theory is true

Scientific change: conjectures and refutations

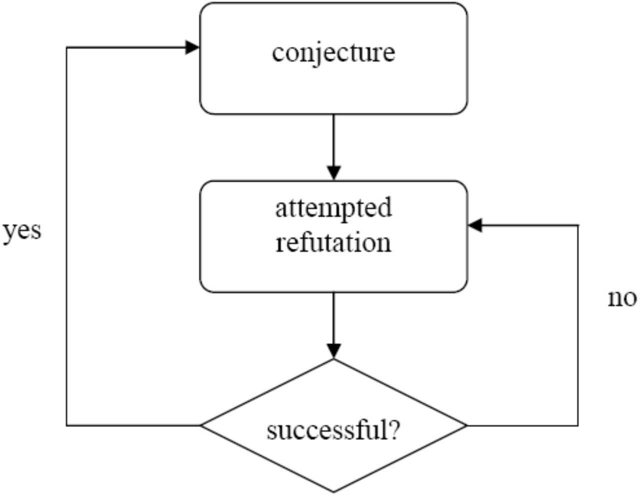

Problems with falsificationism

- holism about testing: no isolated hypothesis can be falsified individually

$\Rightarrow$ which hypothesis ought to be rejected in case of "recalcitrant data"?

- Popper was aware of fact that logic itself does not force a scientist to reject a particular hypothesis in the face of recalcitrant data
- but good scientist would never do that
- any hypothesis can be retained despite apparent falsification if people are only willing to make certain decisions

$\Rightarrow$ scientific thys can be immunized against falsification

Hans Reichenbach (1891-1953)

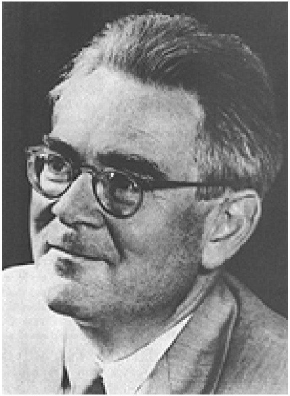

- born in Berlin, emigrated to Turkey in 1933 and to LA in 1938, taught at Berlin, UCLA
- one of the most important representatives of logical empiricism
- founded the Gesellschaft für empirische Philosophie in 1928 in Berlin
- *The Rise of Scientific Philosophy* (1951)

Reichenbach's solution: pragmatic vindication

Earman and Salmon, pp. 64-66.

We don't know whether nature is uniform, so let's draw decision table with two states of the world (Nature is uniform; Nature isn't uniform) and two alternative actions (use induction; don't use induction):

	uniform	not uniform
use induction	success	failure
don't use induction	success or failure	failure

What really requires explanation is lower right-hand box...

Explanation of lower right-hand box

- Suppose crystal gazing were to work consistently

$\Rightarrow$ would be important uniformity that could be found inductively (observations of gazers making successful predictions $\Rightarrow$ crystal gazing will successfully predict future events)

$\Rightarrow$ If crystal gazing can make successful predictions, so can induction.

- But this argument applies to all forms of noninductive reasoning.

$\Rightarrow$ "We therefore have everything to gain and nothing to lose—so far as predicting the future is concerned—by adopting the inductive method." (65)

- Problem: severe vagueness about kind and degree of uniformity needed for the arg to succeed

A brief introduction to confirmation theory

- general goal of confirmation theory: to solve the problem of induction
- more precisely: we have seen that predictions about the future, as well as unrestricted universal generalizations are not logically implied by observational evidence, bc the latter is always only about particular facts in the present and the past
- nevertheless, there is a sense in which observing white swans confirms the hypothesis that the next observed swan is white, and that all swans are white

Characterization (Confirmation theory)

Confirmation theory is the, sometimes formal, attempt to make sense of such confirmation in the wake of the problem of induction.

Models of confirmation of scientific hypotheses

Model (Instantial model of inductive confirmation)

A hypothesis of the form "All F's are G" is supported by its positive instances, i.e. by observed F's that are also G.

(This is sometimes called Nicod confirmation)

Problems:

- observed instances not necessary for inductive support: inference to unobserved entities
- Hempel's paradox of the ravens (to be explained shortly)
- Goodman's "new riddle of induction" (to be explained shortly)

Model (Hypothetico-deductive model of confirmation (Hempel))

A hypothesis or theory is confirmed if it, together with auxiliary statements, deductively entails a datum.

Attractive features:

- allows for confirmation of hypotheses that appeal to unobservable entities and processes, as long as it has observable consequences
- "reduces" inductive inferences to much better understood deductive principles
- seems to genuinely reflect scientific practice, it's "the scientists' philosophy of science" (Lipton, p. 422)

Problems of the hypothetico-deductive model:

1. Hempel's paradox of the ravens
2. Goodman's "new riddle of induction" (curve-fitting problem)

Carl Gustav Hempel (1905-1997): logical empiricism

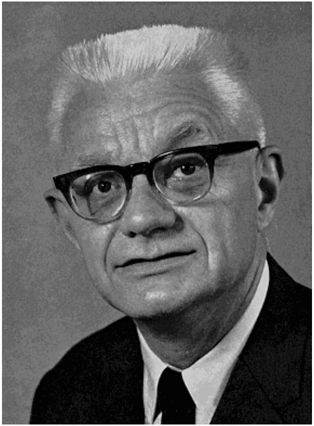

- one of the main representatives of logical empiricism
- studied at Göttingen, Heidelberg, Berlin (PhD 1934)
- 1937 emigration to USA
- taught at Chicago, City College of New York, Yale, Princeton, Pittsburgh
- deductive-nomological model of explanation, hypothetico-deductive model of confirmation

Hempel's raven paradox

Two important principles of confirmation:

1. Equivalence condition: if evidence E confirms hypothesis H_1, and hypothesis H_2 is logically equivalent to H_1, then E also confirms H_2
2. Instance condition: universal generalizations are confirmed by their positive instances

To illustrate the instance condition, consider the universal generalization

H_1: "All ravens are black."

Pedantically, H_1 asserts that: For any x, if x is a raven, then x is black.

Diagrammatically:

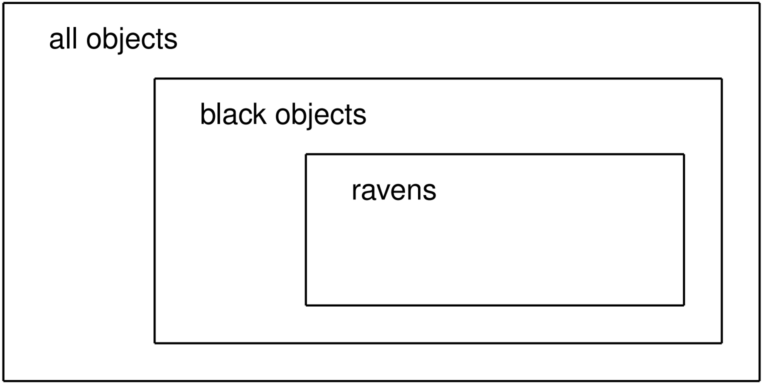

- Let E_1 be the evidence that object a is a raven and that a is black.
- Since the object a satisfies both the antecedent and the consequent of the ravens hypothesis H_1, we have a positive instance of H_1.
- By the instance condition then, E_1 confirms H_1.

Now consider the generalization

H_2: "All non-black things are non-ravens."

Pedantically, H_2 asserts that: For any x, if x is not black, then x is not a raven.

Diagrammatically:

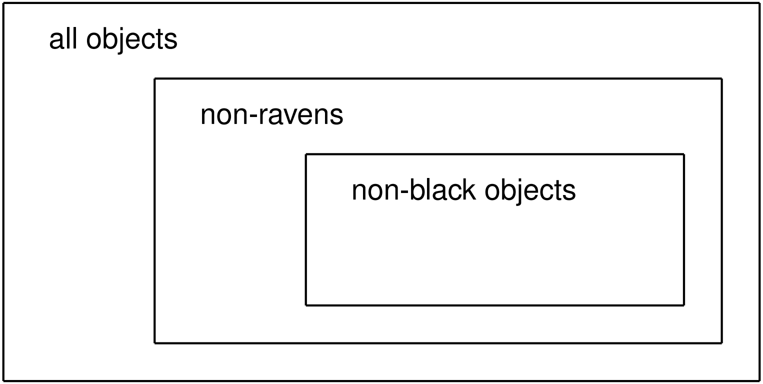

- Let evidence E_2 be the evidence that b is white and that b is a shoe.
- Since b satisfies both the antecedent and the consequent of H_2 we have a positive instance.
- So by the instance condition E_2 confirms H_2.
- But note that H_2 is logically equivalent to H_1.
- So by the equivalence condition, E_2 confirms H_1, i.e. a white shoe confirms "All ravens are black"!
- Does this mean that indoor ornithology is possible?

Resolutions

1. reject equivalence condition not very attractive
2. reject instance condition not very attractive, but we might modify it...
3. H_1 about *ravens*, so E_2 does not really confirm it $\Rightarrow$ test or relevance requirement: objects must be potential falsifiers; ravens are potential falsifiers, but shoes are not
4. swallow consequence:
   - a consider H_3: "All sodium salt burns yellow," but chemical at issue does not burn yellow, and subsequent analysis shows that it's not sodium salt $\Rightarrow$ may count as weak confirmation, although analogous to raven example
   - b in our world, set of non-black things $\gg$ set of ravens; E_2 exhausts a little bit of instances and thereby confirms H_1 a little bit; possible world with ravens $\gg$ non-black objects $\Rightarrow$ more confirmation (Hempel's reply)

But next paradox suggests rejection of instance condition...

Nelson Goodman (1906-1998)

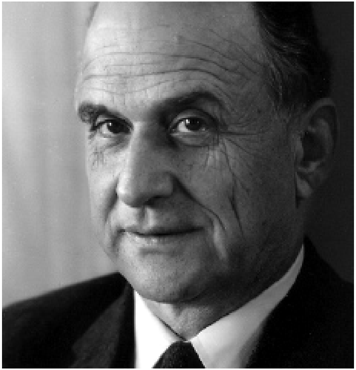

- studied at Harvard (PhD 1941)
- taught at Tufts, U of Pennsylvania, Brandeis, Harvard (his students include Noam Chomsky and Hilary Putnam)
- contributions in aesthetics, epistemology, philosophy of science, and philosophy of language
- was "at odds with rationalism and empiricism alike, with materialism and idealism and dualism, with essentialism and existentialism, with mechanism and vitalism, with mysticism and scientism, and with most other ardent doctrines."

Goodman's "new riddle of induction"

Consider the following argument:

(E_1) *raven* a_1 & *black* a_1
(E_2) *raven* a_2 & *black* a_2
...
$(E_{10,000})$ *raven* $a_{10,000}$ & *black* $a_{10,000}$

(H_1) *All ravens are black.*

Now consider the alternative argument:

(E_1) *raven* a_1 & *blite* a_1
(E_2) *raven* a_2 & *blite* a_2
...
$(E_{10,000})$ *raven* $a_{10,000}$ & *blite* $a_{10,000}$

(H_4) *All ravens are blite.*

Gruesome predicates

The second argument used a new predicate:

Definition (Blite)

An object is blite iff it was first observed before 2010CE and is black, or if it was not first observed before 2010CE and is white.

⚠ Objects do not have to change colour in order to be blite!

If all evidence E_1 through $E_{10,000}$ is based on observation made before 2010CE, then the second argument should be considered as good as the first...

Resolutions

1. reject instance condition
2. only allow "projectable" predicates, i.e. ones not needing a reference to a particular time, or ones that are parasitic on other predicates (black and white in this case)
3. base predicates in language on "natural kinds"

Problem with second resolution:

Definition (Whack)

An object is whack iff it was first observed before 2010CE and is white, or if it was not first observed before 2010CE and is black.

Now consider blite and whack as basic and black and white as parasitic...

Definition (Black)

An object is black iff it was first observed before 2010CE and is blite, or if it was not first observed before 2010CE and is whack.

An unsettling conclusion...

Goodman's new riddle of induction shows that it's actually much worse than Hume thought:

- Hume's solution to his problem of induction doesn't explain why some forms of constant conjunction ("white", "black") give rise to habits of expectation, whereas others don't ("blite", "whack")...

Application: curve-fitting problem

The problem of alternative hypotheses: Boyle's Law

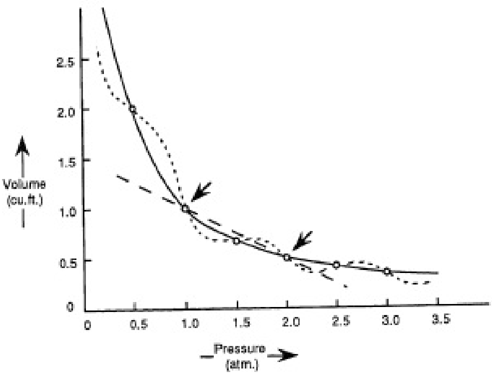

Figure: Boyle's Law (solid line) and alternative laws (from Earman and Salmon, p. 48)

⇒ There's always an infinity of mutually contradictory hypotheses that fit the data, but which is best confirmed?

Other approaches to confirmation

- Carnap's application of the mathematical theory of probability and its present-day successor theory of Bayesianism
- models of causal inference (from effects to their probable causes), such as Mill's "methods of experimental enquiry"
- learning theory
- ...

Bayesian epistemology/Bayesian confirmation theory

- one of the most important developments in epistemology of C20
- offers a formal and mathematically rigorous framework of relating beliefs in hypotheses and evidence confirming or disconfirming it
- framework is probabilistic: assigns probabilities to beliefs
- General idea: a piece of evidence e confirms a hypothesis h in case it raises the probability of h
- probabilities should be "updated" in a way predicted by Bayes's theorem, such that updated degree of belief in hypothesis is probability of hypothesis conditional on evidence
- includes a pragmatic "self-defeat test" for epistemic rationality (next best thing to justification based on deductive logic)

$\Rightarrow$ laws of probability calculus as constraints on rational degrees of belief (or of confidence)

Andrey Nikolaevich Kolmogorov (1903-1987)

- Russian mathematician, Moscow State U
- contributions in probability thy, topology, intuitionistic logic, turbulence, classical mechanics, computational complexity
- main accomplishment: axiomatic foundation of probability theory
- “The theory of probability as mathematical discipline can and should be developed from axioms in exactly the same way as geometry and algebra.”

Kolmogorov axioms of probability theory

Given: class $\mathcal{S}$ of propositions $a, b, c, \ldots$

Introduce probability function on $\mathcal{S}$ as map from $\mathcal{S}$ to the closed interval $[0, 1]$ such that the following axioms hold:

Axiom (1: Non-negativity)

$P(x) \geq 0$ for all x in $\mathcal{S}$; i.e. all probabilities are non-negative.

Axiom (2: Unit measure)

$P(x) = 1$ if x in $\mathcal{S}$ is a tautology; i.e. if x is a proposition that is true in all possible cases, then it has a probability of 1.

Axiom (3: Additivity)

For all x and y in $\mathcal{S}$, if x and y are mutually exclusive propositions, then $P(x \vee y) = P(x) + P(y)$.

Conditional probability

Definition (Conditional probability)

Given a probability function $P(x)$ as defined on the previous slide, the conditional probability $P(s|t)$ of s given t is defined as

$$P(s|t) \doteq \frac{P(s\&t)}{P(t)}$$

Bayesian updating of beliefs

Bayesians make (a more complicated version of) the following epistemological assumption:

Principle (Conditionalization)

Starting from initial (prior) probabilities $P_i(h)$ of any hypothesis statement h, acquiring new evidence in the sense of becoming certain of the evidence statement e, rationality dictates that one updates one's initial probabilities to obtain one's final (posterior) probabilities by "conditionalizing" on e

$$P_i(h) \longrightarrow P_f(h) = P_i(h|e)$$

So we should find a way to calculate $P_i(h|e)$ as a function of $P_i(h)$.

Theorem by Rev Thomas Bayes (1702-1761)

For all propositions h and e, we have

$$P(h|e) = \frac{P(h) \cdot P(e|h)}{P(e)} = \frac{P(h) \cdot P(e|h)}{P(e|h) \cdot P(h) + P(e|\neg h) \cdot P(\neg h)}$$

where

- $P(h)$: prior probability of h
- $P(h|e)$: posterior probability of h (in the light of e)
- $P(e|h)$: "likelihood" of evidence e on hypothesis h

Bayesian updating

- first, determine the prior probability of h and the likelihood that e_1 will be observed given h
- determine the probability to observe e_1 *independently of h*
- if e_1 is observed, calculate the posterior probability $P(h|e_1)$ via Bayes's theorem
- consider this posterior probability as your new prior probability of h
- consider the probability of a new piece of evidence e_2 and its likelihood in the light of h
- if e_2 is observed, calculate the new posterior probability of h via Bayes's theorem
- ...

Example 1: from which bowl is the cookie?

Two bowls of cookies:

1. $Bowl_1$ has 10 chocolate chip and 30 plain cookies
2. $Bowl_2$ has 20 chocolate chip and 20 plain cookies

Question: If you pick a random cookie from a random bowl, and it is plain (e), how probable is it that it's from $Bowl_1$ (h)?

Priors: $P(h) = P(\neg h) = 0.5$

Likelihoods: $P(e|h) = 0.75$ and $P(e|\neg h) = 0.5$

Use Bayes's theorem:

$$\begin{aligned} P(h|e) &= \frac{P(h) \cdot P(e|h)}{P(e|h) \cdot P(h) + P(e|\neg h) \cdot P(\neg h)} \\ &= \frac{0.5 \cdot 0.75}{0.5 \cdot 0.75 + 0.5 \cdot 0.5} = 0.6. \end{aligned}$$

Example 2: is she/he going to the party?

from Peter Godfrey-Smith, *Theory and Reality: An Introduction to the Philosophy of Science*, p. 204.

- h: hypothesis that she/he is at party
- e: evidence that her/his car is parked outside
- $P(h)$: initial probability that she/he is at party (before seeing the car); let's say this is 0.5.
- $P(e|h)$: likelihood that her/his car is parked outside if she/he is at the party; suppose this is 0.8
- $P(e|\neg h)$: likelihood that her/his car is parked outside if she/he is not at the party; suppose this is only 0.1
- $P(h|e)$: prob that she/he is at the party given that her/his car is parked outside; can be calculated using Bayes's theorem:

$$P(h|e) = \frac{0.5 \cdot 0.8}{0.5 \cdot 0.8 + 0.5 \cdot 0.1} = 0.89.$$

$\Rightarrow$ seeing the car raises the prob of h from 0.5 to 0.89.

Example 3: in the courtroom

Juror must assess how evidence bears on guilt of defendant:

- g: hypothesis that defendant is guilty
- e: evidence that defendant's DNA matches DNA found at crime scene
- $P(e|g)$: likelihood to see evidence of matching DNA if defendant is guilty; in capital offenses, typically very high; here assumed to be 1
- $P(e|\neg g)$: likelihood to see evidence of matching DNA if defendant is not guilty; very low, assume 1 in a million, or 10^{-6}
- $P(g)$: initial probability that defendant is guilty (prior); hugely depends on other evidence, circumstances etc. Two cases: either (A) strong prior suspicion ($P(g) = 0.3$), or (B) very low suspicion ($P(g) = 10^{-6}$)
- $P(g|e)$: prob that defendant is guilty if matching DNA is found; this is what we want to know!

Case (A): $P(g) = 0.3$

$$\begin{aligned}
P(g|e) &= \frac{P(g)\cdot P(e|g)}{P(e|g)\cdot P(g) + P(e|\neg g)\cdot P(\neg g)} \\
&= \frac{0.3 \cdot 1.0}{0.3\cdot 1.0 + 0.7 \cdot 10^{-6}} \\
&= 0.99999766667
\end{aligned}$$

Case (B): $P(g) = 10^{-6}$

$$\begin{aligned}
P(g|e) &= \frac{P(g)\cdot P(e|g)}{P(e|g)\cdot P(g) + P(e|\neg g)\cdot P(\neg g)} \\
&= \frac{10^{-6} \cdot 1.0}{10^{-6}\cdot 1.0 + (1-10^{-6}) \cdot 10^{-6}} \\
&\approx 0.5
\end{aligned}$$

Example 4: the *Scorpion* search

- May 1968: US nuclear submarine *Scorpion* fails to arrive at home port of Norfolk, VA
- US Navy: convinced that vessel had been lost off Eastern seabord, but extensive search fails to discover wreck
- US Navy deep water expert John Craven believed that it was southwest of Portuguese archipelago of the *Azores* based on controversial triangulation of hydrophones
- allocation of limited resources (one ship) $\Rightarrow$ optimize them
- Craven worked with mathematicians to optimize the search, using Bayesian search theory
- October 1968: wreck is found 400 miles southwest of the Azores

Bayesian search theory

1. Sea is divided into grid squares
2. experienced submarine commanders are interviewed, etc, to formulate a number of hypotheses about what happened to vessel
3. construct prob distribution over squares corresponding to each hypothesis
4. construct prob distribution for actually finding object in square X if it really is in X (function of water depth)
5. combine all these prob distributions (from 3 and 4) to produce overall probability grid; this gives prob of finding object in a square if this square is searched (for all squares)
6. construct a search path starting from square of highest prob that then searches high prob areas, then intermediate prob areas, then low prob areas
7. revise overall prob distribution continuously as you search, i.e. if you have unsuccessfully searched square, then prob that object is there is greatly reduced (though usually not zero), and prob of finding it elsewhere must be increased; this revision follows Bayes's theorem

Example 5: V1@gra and Bayesian filtering

`Introduction to Bayesian filtering`

- task: figure out how likely an email message is spam based on the words that appear in it
- basic idea: Bayesian spam filter
    - lists words in incoming emails,
    - assigns to each word probability that it appears in spam mail (misspelled words score very high), and
    - uses these probs as input into Bayes's formula to determine whether or not email is spam
- first need to train the spam filter by showing it spam and non-spam mails (more and more automated)
- spam filter stores all words in trained messages (incl host name, IP address, HTML tag, etc) in databases

$\Rightarrow$ filter calculates prob how likely it is that a word is in spam based on its frequency in databases (the “spamicity” of each word)

- spamicity of 0.5 is neutral, higher (lower) means that it often occurs in (non-)spam messages
- filter then uses Bayes’s formula to calculate the overall spamicity of a message based on the spamicity of all the words that occur in it

$\Rightarrow$ message put in spam filter if spamicity is above 0.5

- Generally: Bayesian spam filter are highly effective bc (1) they adapt to individual circumstances (databases are built for each user), and (2) they learn over time and update the databases

Example 6: Predicting the US presidential election

`How many electoral votes will Obama get?`

- The link above leads to a blog where a simple Bayesian method is used to predict the outcome of the last US presidential election (on 31 October 2008).
- The blogger predicted, based on polling data from cnn.com and 5000 simulated elections, that the probability that McCain would win enough electoral votes to win him the White House was 0.0.

What is probability?

(objective) chance	vs.	(subjective) degree of belief
↓		↓
How to measure?		How to measure?
↓		↓
frequency of occurrence		gambling behaviour (via wagers)

Axioms of probability must apply to both!

Subjectivist Bayesianism

- "probability as degree of personal belief" (generated by free choice, socialization, evolution, etc)
- The probability of an event is just the certainty with which a Bayesian agent expects the event to occur.
- Main idea: "rational belief" should be understood as a generalization of betting behaviour: given an amount of information/data and asked for an evaluation, what odds would one bet for the truth of one's evaluation?
- to bet on h at odds of $X : 1$ is to be willing to risk losing \$$X$ if h is false, in return for a gain of \$1 if h is true
- if your subjectively fair odds for a bet on h are $X : 1$, then your degree of belief in h is $X/(X+1)$
- It is of course possible that the subjective degree of belief violates the axioms, but...

"Dutch book" theorem

If somebody's subjective degree of belief violates a Kolmogorov axiom, then this person should accept a combination of bets which amounts to a so-called Dutch book, i.e.

the combination of bets that they should accept guarantees the person a loss!

Simple example: your degree of belief that next coin toss will come out "heads" is 0.55, and your degree of belief that it'll be "tails" is 0.5 $\Rightarrow$ bookie can write out set of wagers which guarantee that you'll lose 5c on each dollar you bet

Bayesian solution to grue paradox

- Suppose you're presented two inductive arguments from the same set of observations of green emeralds, one arguing that all emeralds are green, the other that they are grue.
- Why is one induction better than the other?
- Standard Bayesian answer: both are OK, but most people would assign higher prior prob to "green" hypothesis than to "grue" hypothesis
- Reaction: true, it gives a difference, but does it explain why the "grue" hypothesis gets lower prior prob?
- Bayesianism offers no criticism of subjective decision to assign high prior prob to "grue" hypothesis, as long as probs are internally coherent and updated properly

Problem 1: Priors

- initial set of prior probabilities can be chosen freely (except for 0 and 1)
- but how could a strange assignment of priors be criticized, so long as it follows the axioms?
- Bayesian answer: doesn't matter bc initial set of priors are washed out asymptotically (convergence, stable estimation theorem)
- problem: conversely, we also have for any amount of evidence, and any measure of agreement, there is some set of priors s.t. this evidence will *not* get the two people to agree by the end (Kyburg)
- there must be agreement concerning the likelihoods $P(e_i|h)$, relevance of particular pieces of evidence
- assumptions of theorems do not even remotely apply in realistic scientific contexts

Problem 2: Old evidence (Glymour)

- problem of old evidence: old evidence can in fact confirm new thy, but according to Bayesian kinematics it cannot
- suppose e is known before theory T is introduced at time t
- bc e is known at t, $P_t(e) = 1$

$\Rightarrow$ likelihood of e given T is also 1: $P_t(e|T) = 1$

$$P_t(T|e) = \frac{P_t(T) \cdot P_t(e|T)}{P_t(e)} = P_t(T)$$

$\Rightarrow$ posterior prob of T is same as its prior prob!

By way of conclusion

Hájek and Hartmann, “Bayesian epistemology”, *A Companion to Epistemology*, Oxford: Blackwell, 2009.

Alan Hájek and Stephan Hartmann contrast two views on Bayesian epistemology:

> *According to one view, there cannot [be a Bayesian epistemology]: Bayesianism fails to do justice to essential aspects of knowledge and belief, and as such it cannot provide a genuine epistemology at all. According to another view, Bayesianism should supersede traditional epistemology: where the latter has been mired in endless debates over skepticism and Gettierology, Bayesiansim offers the epistemologist a research program. We will advocate a more moderate view: Bayesianism can illuminate various long-standing problems of epistemology, while not addressing all of them; and while Bayesianism opens up fascinating new areas of research, it by no means closes down the staple preoccupations of traditional epistemology. (93)*