Confirmation Theory

J. Dmitri Gallow

July 7th, 2017

Pittsburgh Summer Program
University of Pittsburgh

Please interrupt

Induction & Deduction

- In a *deductive* inference, the truth of the premises *guarantees* the truth of the conclusion.

Induction & Deduction

- In a *deductive* inference, the truth of the premises *guarantees* the truth of the conclusion.

P1. If H, then E

Induction & Deduction

- In a *deductive* inference, the truth of the premises *guarantees* the truth of the conclusion.

P1. If H, then E

P2. H

Induction & Deduction

- In a *deductive* inference, the truth of the premises *guarantees* the truth of the conclusion.

P1. If H, then E

P2. H

C. E

Induction & Deduction

- In an *inductive* inference, the truth of the premises *doesn't* guarantee the truth of the conclusion.

Induction & Deduction

- In an *ampliative* inference, the truth of the premises *doesn't* guarantee the truth of the conclusion.

P1. If H, then E

Induction & Deduction

- In an *inductive* inference, the truth of the premises *doesn't* guarantee the truth of the conclusion.

P1. If H, then E

P2. E

Induction & Deduction

- In an *inductive* inference, the truth of the premises *doesn't* guarantee the truth of the conclusion.

P1. If H, then E

P2. E

C. H

Deductive Logic

- Deductive Logic is the study of which deductive inferences are good *qua* deductive inferences (valid) and which are bad *qua* deductive inferences (invalid).

Deductive Logic

- Deductive Logic is the study of which deductive inferences are good *qua* deductive inferences (valid) and which are bad *qua* deductive inferences (invalid).

$$\begin{array}{l} H \to E \\ H \\ \hline E \end{array} \qquad\qquad \begin{array}{l} H \to E \\ \neg H \\ \hline \neg E \end{array}$$

Inductive Logic & the Theory of Confirmation

- Inductive Logic is the study of which inductive inferences are good *qua* inductive inferences and which are bad *qua* inductive inferences.

Inductive Logic & the Theory of Confirmation

- Confirmation Theory is the study of which inductive inferences are good *qua* inductive inferences and which are bad *qua* inductive inferences.

Inductive Logic & the Theory of Confirmation

- Inductive Logic is the study of which inductive inferences are good *qua* inductive inferences and which are bad *qua* inductive inferences.

Inductive Logic & the Theory of Confirmation

- Inductive Logic is the study of which inductive inferences are good *qua* inductive inferences and which are bad *qua* inductive inferences.

$$\frac{\text{All observed swans are white.}}{\text{All swans are white.}}$$

Inductive Logic & the Theory of Confirmation

- Inductive Logic is the study of which inductive inferences are good *qua* inductive inferences and which are bad *qua* inductive inferences.

$$\frac{\text{All observed swans are white.}}{\text{All swans are white.}}$$

$$\frac{\text{All observed swans are white.}}{\text{All but the observed swans are black.}}$$

Two Questions About Deduction

- Question 1: Which deductive inferences are good, and which are bad? (What are the canons of deductive logic?)

Two Questions About Deduction

- Question 1: Which deductive inferences are good, and which are bad? (What are the canons of deductive logic?)
- Question 2: Why should we think that good deductive inferences will lead us to truth?

Two Questions About Induction

- Question 1: When is an inductive inference good, and when it is bad? (What are the canons of inductive logic/the theory of confirmation?)

Two Questions About Induction

- Question 1: When is an inductive inference good, and when it is bad? (What are the canons of inductive logic/the theory of confirmation?)
- Question 2: Why should we think that good inductive inferences will lead us to truth?

Two Questions About Induction

- Question 1: When is an inductive inference good, and when it is bad? (What are the canons of inductive logic/the theory of confirmation?)
- Question 2: Why should we think that good inductive inferences will lead us to truth?

▷ David Hume: there is no non-circular answer to Question 2.

Table of contents

Confirmation & Disconfirmation

Confirmation & Disconfirmation

- Sometimes, a piece of evidence, *E*, gives reason to believe a hypothesis, *H*.
    - When this is so, say that *E confirms H*.
- Other times, a piece of evidence, *E*, gives reason to *dis*believe a hypothesis, *H*.
    - When this is so, say that *E disconfirms H*.

Confirmation & Disconfirmation

- Sometimes, a piece of evidence, *E*, gives reason to believe a hypothesis, *H*.
  - When this is so, say that *E confirms H*.
- Other times, a piece of evidence, *E*, gives reason to *dis*believe a hypothesis, *H*.
  - When this is so, say that *E disconfirms H*.

Confirmation & Disconfirmation

- Sometimes, a piece of evidence, *E*, gives reason to believe a hypothesis, *H*.
  - When this is so, say that *E confirms H*.
- Other times, a piece of evidence, *E*, gives reason to *dis*believe a hypothesis, *H*.
  - When this is so, say that *E disconfirms H*.

Confirmation & Disconfirmation

- Sometimes, a piece of evidence, *E*, gives reason to believe a hypothesis, *H*.
    - When this is so, say that *E confirms H*.
- Other times, a piece of evidence, *E*, gives reason to *dis*believe a hypothesis, *H*.
    - When this is so, say that *E disconfirms H*.

Confirmation *vs* Belief

- Just because we have some evidence, E, which confirms H, this doesn't mean that we should think H is true.

Confirmation is *Degreed*

- Confirmation comes in degrees.
- *E* could give a very *strong* reason to believe that *H*, or it could give a rather *weak* reason to believe that *H*.
- If *E* confirms *H*, but only very weakly, then it could be that we shouldn't believe *H*.

Confirmation is *Degreed*

- Confirmation comes in degrees.
- E could give a very *strong* reason to believe that H, or it could give a rather *weak* reason to believe that H.
- If E confirms H, but only very weakly, then it could be that we shouldn't believe H.

Confirmation is *Degreed*

- Confirmation comes in degrees.
- *E* could give a very *strong* reason to believe that *H*, or it could give a rather *weak* reason to believe that *H*.
- If *E* confirms *H*, but only very weakly, then it could be that we shouldn't believe *H*.

Confirmation is *Degreed*

Confirmation *vs* Belief

- Just because we have some evidence, E, which confirms H, this doesn't mean that we should think H is true.

Evidential Defeat

- Let H be the hypothesis that John robbed the bank, and let E be the evidence that 100 eyewitnesses who know John personally identified John as the bank robber.
- E strongly confirms H. But just because you have the evidence E, this doesn't mean that you should believe H.
- E needn't be your *total* evidence.
- You could additionally have the evidence that John has an identical twin brother and that John has a rock-solid alibi.

Evidential Defeat

- Let H be the hypothesis that John robbed the bank, and let E be the evidence that 100 eyewitnesses who know John personally identified John as the bank robber.
- E strongly confirms H. But just because you have the evidence E, this doesn't mean that you should believe H.
- E needn't be your *total* evidence.
- You could additionally have the evidence that John has an identical twin brother and that John has a rock-solid alibi.

Evidential Defeat

- Let H be the hypothesis that John robbed the bank, and let E be the evidence that 100 eyewitnesses who know John personally identified John as the bank robber.
- E strongly confirms H. But just because you have the evidence E, this doesn't mean that you should believe H.
- E needn't be your *total* evidence.
- You could additionally have the evidence that John has an identical twin brother and that John has a rock-solid alibi.

Evidential Defeat

- Let H be the hypothesis that John robbed the bank, and let E be the evidence that 100 eyewitnesses who know John personally identified John as the bank robber.
- E strongly confirms H. But just because you have the evidence E, this doesn't mean that you should believe H.
- E needn't be your *total* evidence.
- You could additionally have the evidence that John has an identical twin brother and that John has a rock-solid alibi.

Evidential Defeat

- Deductive inference is monotonic (or indefeasible)
  - If *P deductively entails C*, then *P&Q* deductively entails *C* as well.
- Inductive inference is non-monotonic (or defeasible)
  - If *E* confirms *H*, it doesn't follow that *E&F* confirms *H*.

Evidential Defeat

- Deductive inference is monotonic (or indefeasible)
    - If P *deductively entails* C, then $P\&Q$ deductively entails C as well.
- Inductive inference is non-monotonic (or defeasible)
    - If E confirms H, it doesn't follow that $E\&F$ confirms H.

Evidential Defeat

- Deductive inference is monotonic (or indefeasible)
    - If P *deductively entails* C, then $P\&Q$ deductively entails C as well.
- Inductive inference is non-monotonic (or defeasible)
    - If E confirms H, it doesn't follow that $E\&F$ confirms H.

Evidential Defeat

- Deductive inference is monotonic (or indefeasible)
  - If P *deductively entails* C, then $P\&Q$ deductively entails C as well.
- Inductive inference is non-monotonic (or defeasible)
  - If E confirms H, it doesn't follow that $E\&F$ confirms H.

What We Want from a Theory of Confirmation

- A *qualitative* account of confirmation.
  - For any H, E: does E confirm H?
- A *quantitative* measure of confirmation.
  - For any H, E: to what *degree* does E confirm H?
- We'd like our theory of confirmation to be *formal* and *intersubjective*.
  - *Formal*: we can say whether E confirms H by looking only at syntax, or logical *form*.
  - *Intersubjective*: we can all agree about whether E confirms H

What We Want from a Theory of Confirmation

- A *qualitative* account of confirmation.
    - For any H, E: does E confirm H?
- A *quantitative* measure of confirmation.
    - For any H, E: to what *degree* does E confirm H?
- We'd like our theory of confirmation to be *formal* and *intersubjective*.
    - *Formal*: we can say whether E confirms H by looking only at syntax, or logical *form*.
    - *Intersubjective*: we can all agree about whether E confirms H

What We Want from a Theory of Confirmation

- A *qualitative* account of confirmation.
    - For any H, E: does E confirm H?
- A *quantitative* measure of confirmation.
    - For any H, E: to what *degree* does E confirm H?
- We'd like our theory of confirmation to be *formal* and *intersubjective*.
    - *Formal*: we can say whether E confirms H by looking only at syntax, or logical *form*.
    - *Intersubjective*: we can all agree about whether E confirms H

What We Want from a Theory of Confirmation

- A *qualitative* account of confirmation.
    - For any H, E: does E confirm H?
- A *quantitative* measure of confirmation.
    - For any H, E: to what *degree* does E confirm H?
- We'd like our theory of confirmation to be *formal* and *intersubjective*.
    - *Formal*: we can say whether E confirms H by looking only at syntax, or logical *form*.
    - *Intersubjective*: we can all agree about whether E confirms H

What We Want from a Theory of Confirmation

- A *qualitative* account of confirmation.
    - For any H, E: does E confirm H?
- A *quantitative* measure of confirmation.
    - For any H, E: to what *degree* does E confirm H?
- We'd like our theory of confirmation to be *formal* and *intersubjective*.
    - *Formal*: we can say whether E confirms H by looking only at syntax, or logical *form*.
    - *Intersubjective*: we can all agree about whether E confirms H

What We Want from a Theory of Confirmation

- A *qualitative* account of confirmation.
    - For any H, E: does E confirm H?
- A *quantitative* measure of confirmation.
    - For any H, E: to what *degree* does E confirm H?
- We'd like our theory of confirmation to be *formal* and *intersubjective*.
    - *Formal*: we can say whether E confirms H by looking only at syntax, or logical *form*.
    - *Intersubjective*: we can all agree about whether E confirms H

What We Want from a Theory of Confirmation

- A *qualitative* account of confirmation.
    - For any H, E: does E confirm H?
- A *quantitative* measure of confirmation.
    - For any H, E: to what *degree* does E confirm H?
- We'd like our theory of confirmation to be *formal* and *intersubjective*.
    - *Formal*: we can say whether E confirms H by looking only at syntax, or logical *form*.
    - *Intersubjective*: we can all agree about whether E confirms H

You Can't Always Get What You Want

Hempel's Impossibility Results

- A promising first thought: deductive consequences of a hypothesis confirm it.

Entailments Confirm

If *H entails E*, then *E confirms H*.

Hempel's Impossibility Results

- A promising first thought: deductive consequences of a hypothesis confirm it.

Entailments Confirm

If *H entails E*, then *E confirms H*.

Hempel's Impossibility Results

- Another promising thought: confirmation transmits through deduction.

Consequence Condition

If E confirms H, then E confirms anything which H entails.

Hempel's Impossibility Results

- Another promising thought: confirmation transmits through deduction.

Consequence Condition

If E confirms H, then E confirms anything which H entails.

Hempel's Impossibility Results

Entailments Confirm

If *H entails E*, then *E confirms H*.

Consequence Condition

If *E* confirms *H*, then *E* confirms anything which *H* entails.

- If we accept both Entailments Confirm and the Consequence Condition, then we must say that every proposition confirms every other proposition.

Hempel's Impossibility Results

$$\frac{A \mathbin{\&} B}{A}$$

Hempel’s Impossibility Results

$$\frac{A \mathbin{\&} B}{A} \qquad \xRightarrow{EC} \qquad \frac{A}{A \mathbin{\&} B}$$

Hempel's Impossibility Results

$$\frac{A \mathbin{\&} B}{A} \qquad \xRightarrow{EC} \qquad \frac{A}{A \mathbin{\&} B}$$

$$\frac{A \mathbin{\&} B}{B}$$

Hempel's Impossibility Results

$$
\begin{tikzcd}
\begin{array}{c} A \,\&\, B \\ \hline A \end{array} \arrow[r, Rightarrow, "EC"] & \begin{array}{c} A \\ \hline\hline A \,\&\, B \end{array} \arrow[d, Rightarrow, "CC"'] \\
\begin{array}{c} A \,\&\, B \\ \hline B \end{array} \arrow[r, Rightarrow] & \begin{array}{c} A \\ \hline\hline B \end{array}
\end{tikzcd}
$$

Hempel's Impossibility Results

Entailments Confirm

If *H entails E*, then *E confirms H*.

Consequence Condition

If *E* confirms *H*, then *E* confirms anything which *H* entails.

- Perhaps we should weaken these principles.

Hempel's Impossibility Results

Laws are Confirmed by Their Instances

A law statement of the form "All Fs are Gs" is confirmed by an F G.

Consequence Condition

If E confirms H, then E confirms anything which H entails.

- Perhaps we should weaken these principles.

Hempel's Impossibility Results

Laws are Confirmed by Their Instances

A law statement of the form "All Fs are Gs" is confirmed by an F G.

Equivalence Condition

If E confirms H, then E confirms anything equivalent to H.

- Perhaps we should weaken these principles.

Hempel's Impossibility Results

Laws are Confirmed by Their Instances

A law statement of the form "All Fs are Gs" is confirmed by an F G.

Equivalence Condition

If E confirms H, then E confirms anything equivalent to H.

- New Problem: *nearly everything* confirms any given law statement.

Hempel's Impossibility Results

1. "All ravens are black" is equivalent to "All non-black things are non-ravens".
2. By Laws are Confirmed by Their Instances, a green leaf (which is both a non-black thing and a non-raven) confirms the hypothesis that all non-black things are non-ravens.
3. By 1, 2, and the Equivalence Condition, a green leaf confirms the hypothesis that all ravens are black.

Hempel's Impossibility Results

1. "All ravens are black" is equivalent to "All non-black things are non-ravens".
2. By Laws are Confirmed by Their Instances, a green leaf (which is both a non-black thing and a non-raven) confirms the hypothesis that all non-black things are non-ravens.
3. By 1, 2, and the Equivalence Condition, a green leaf confirms the hypothesis that all ravens are black.

Hempel's Impossibility Results

1. "All ravens are black" is equivalent to "All non-black things are non-ravens".
2. By Laws are Confirmed by Their Instances, a green leaf (which is both a non-black thing and a non-raven) confirms the hypothesis that all non-black things are non-ravens.
3. By 1, 2, and the Equivalence Condition, a green leaf confirms the hypothesis that all ravens are black.

Goodman's Impossibility Result

- Goodman: there's a deeper problem here. No theory of confirmation can be purely formal.
- In order to say whether a hypothesis of the form "All *F*s are *G*s" is confirmed by an *F G*, we must know something about what '*F*' and '*G*' *mean*.

Goodman's Impossibility Result

- Goodman: there's a deeper problem here. No theory of confirmation can be purely formal.
- In order to say whether a hypothesis of the form "All *F*s are *G*s" is confirmed by an *F G*, we must know something about what '*F*' and '*G*' *mean*.

Goodman's Impossibility Result

The first observed emerald is green
The second observed emerald is green
⋮
The *n*th observed emerald is green

All unobserved emeralds are green

Goodman's Impossibility Result

- Say that a thing is <u>grue</u> iff it has been observed before 2018 and is green or has not been observed before 2018 and is blue.

Goodman's Impossibility Result

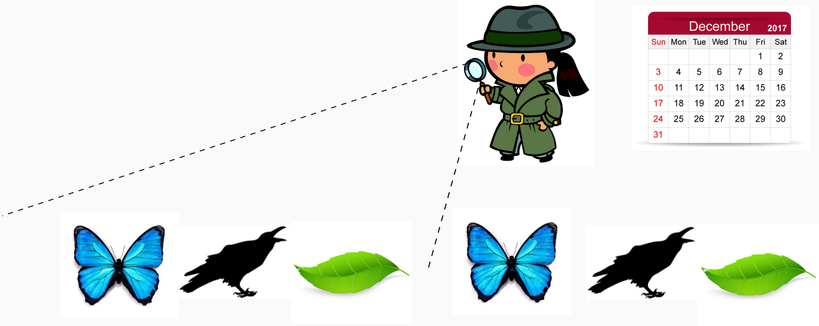

Goodman’s Impossibility Result

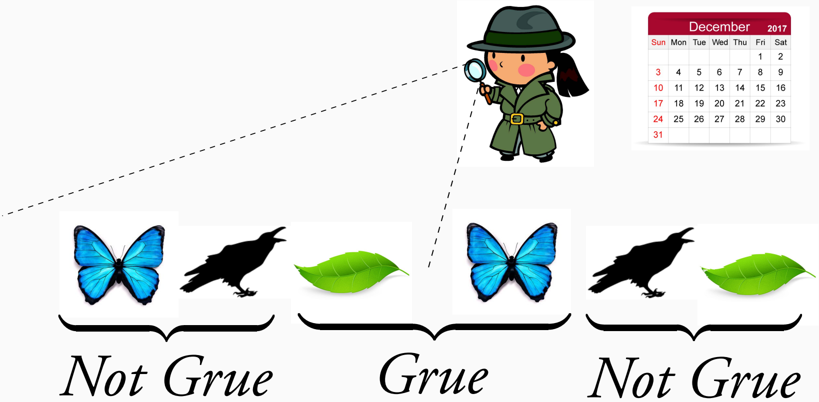

Goodman's Impossibility Result

The first observed emerald is green

The second observed emerald is green

$\vdots$

The *n*th observed emerald is green

$$\overline{\overline{}}$$

All unobserved emeralds are green

Goodman's Impossibility Result

The first observed emerald is grue

The second observed emerald is grue

⋮

The *n*th observed emerald is grue

All unobserved emeralds are green

Goodman's Impossibility Result

The first observed emerald is grue

The second observed emerald is grue

⋮

The *n*th observed emerald is grue

All unobserved emeralds are grue

Goodman's Impossibility Result

$$\dfrac{\begin{array}{l}\text{The first observed emerald is grue}\\ \text{The second observed emerald is grue}\\ \quad\vdots\\ \text{The } n\text{th observed emerald is grue}\end{array}}{\text{All unobserved emeralds are blue}}$$

Goodman's Impossibility Result

- If "All unobserved emeralds are green" is confirmed by the observation of n green emeralds, then so too is the hypothesis that "All unobserved emeralds are blue".
- A purely *formal* theory of confirmation cannot distinguish induction from counter-induction.
- So a theory of induction must go beyond logical form.

Goodman's Impossibility Result

- If "All unobserved emeralds are green" is confirmed by the observation of *n* green emeralds, then so too is the hypothesis that "All unobserved emeralds are blue".
- A purely *formal* theory of confirmation cannot distinguish induction from counter-induction.
- So a theory of induction must go beyond logical form.

Goodman's Impossibility Result

- If "All unobserved emeralds are green" is confirmed by the observation of n green emeralds, then so too is the hypothesis that "All unobserved emeralds are blue".
- A purely *formal* theory of confirmation cannot distinguish induction from counter-induction.
- So a theory of induction must go beyond logical form.

What We Want from a Theory of Confirmation

- A *qualitative* account of confirmation.
    - For any H, E: does E confirm H?
- A *quantitative* measure of confirmation.
    - For any H, E: to what *degree* does E confirm H?
- We'd like our theory of confirmation to be ~~*formal* and~~ *intersubjective*.
    - ~~*Formal*: we can say whether E confirms H by looking only at syntax, or logical *form*.~~
    - *Intersubjective*: we can all agree about whether E confirms H

Confirmation & Probability

Confirmation & Probability

Probability

Probability

- A *probability function*, Pr, is any function from a set of propositions, $\mathscr{P}$, to the unit interval, $[0, 1]$

$$\mathsf{Pr} : \mathscr{P} \to [0, 1]$$

which also has the following properties:

Ax1. If the proposition $\top$ is necessarily true, then $\mathsf{Pr}(\top) = 1$.

Ax2. If the propositions A and B are inconsistent, then $\mathsf{Pr}(A \vee B) = \mathsf{Pr}(A) + \mathsf{Pr}(B)$.

Probability

- A *probability function*, Pr, is any function from a set of propositions, $\mathscr{P}$, to the unit interval, $[0, 1]$

$$\mathsf{Pr} : \mathscr{P} \to [0, 1]$$

which also has the following properties:

Ax1. If the proposition $\top$ is necessarily true, then $\mathsf{Pr}(\top) = 1$.

Ax2. If the propositions A and B are inconsistent, then $\mathsf{Pr}(A \vee B) = \mathsf{Pr}(A) + \mathsf{Pr}(B)$.

Probability

- A *probability function*, Pr, is any function from a set of propositions, $\mathscr{P}$, to the unit interval, $[0, 1]$

$$\mathsf{Pr} : \mathscr{P} \to [0, 1]$$

which also has the following properties:

Ax1. If the proposition $\top$ is necessarily true, then $\mathsf{Pr}(\top) = 1$.

Ax2. If the propositions A and B are inconsistent, then $\mathsf{Pr}(A \vee B) = \mathsf{Pr}(A) + \mathsf{Pr}(B)$.

Probability

- If Pr is a probability function, then we may represent it with a *muddy Venn diagram*.

Probability

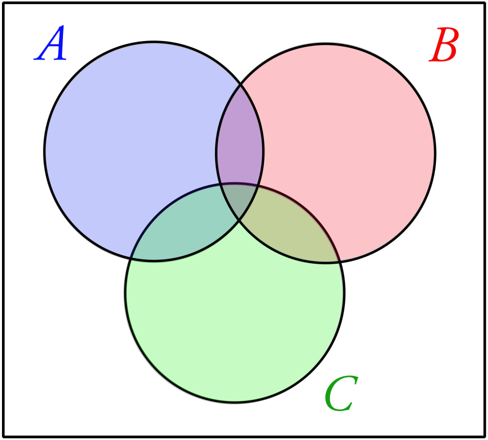

Probability

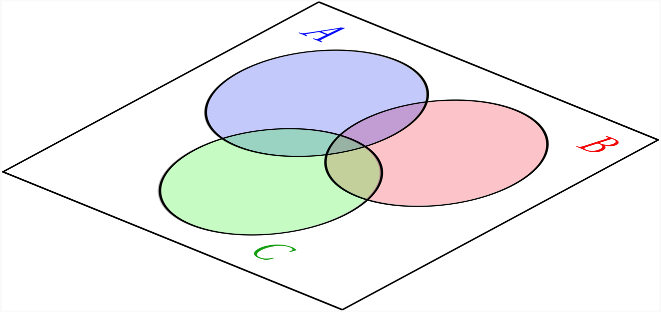

Probability

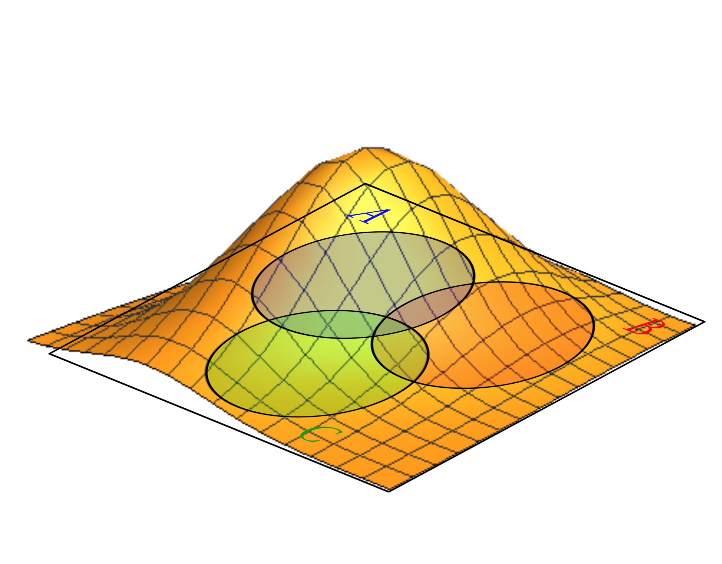

Probability

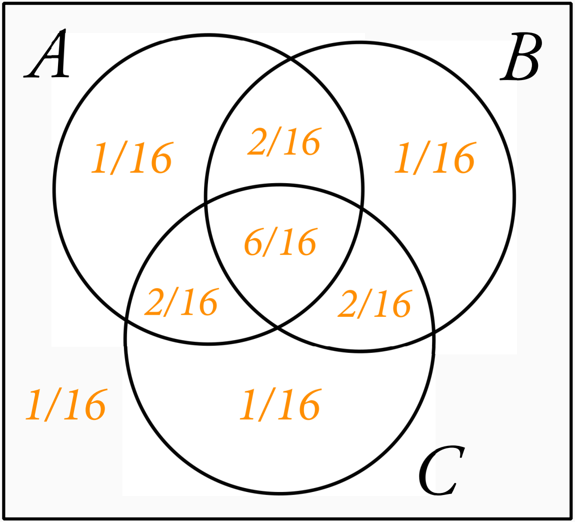

Probability

A	B	C	Pr
T	T	T	6/16
T	T	F	2/16
T	F	T	2/16
T	F	F	1/16
F	T	T	2/16
F	T	F	1/16
F	F	T	1/16
F	F	F	1/16

Conditional Probability

- We introduce the following *definition*:

$$\mathsf{Pr}(A \mid B) \stackrel{\text{def}}{=} \frac{\mathsf{Pr}(A \,\&\, B)}{\mathsf{Pr}(B)}, \quad \text{if defined}$$

Conditional Probability

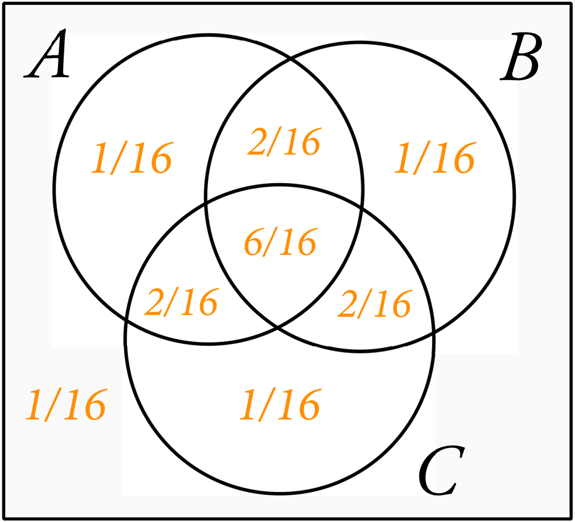

Conditional Probability

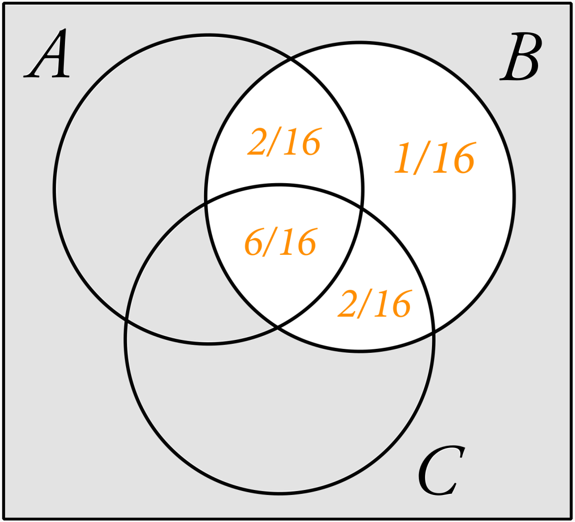

Conditional Probability

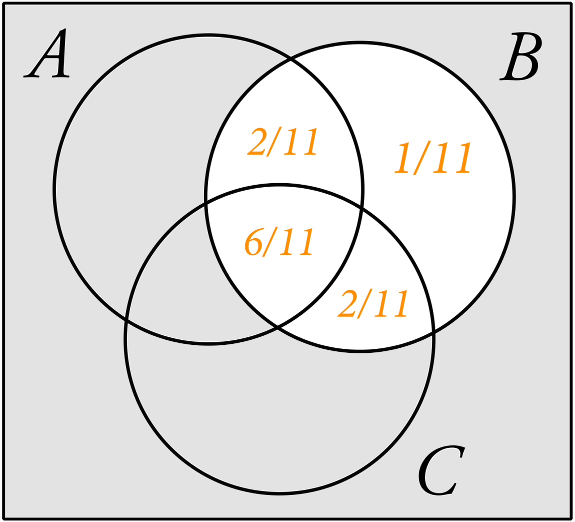

Probabilistic Independence

- We may say that the propositions A and B are *independent* (according to Pr) if and only if

$$\mathsf{Pr}(A \,\&\, B) = \mathsf{Pr}(A) \cdot \mathsf{Pr}(B)$$

Probabilistic Independence

- We may say that the propositions A and B are *independent* (according to Pr) if and only if

$$\frac{\mathsf{Pr}(A \,\&\, B)}{\mathsf{Pr}(B)} = \mathsf{Pr}(A)$$

Probabilistic Independence

- We may say that the propositions A and B are *independent* (according to Pr) if and only if

$$\mathsf{Pr}(A \mid B) = \mathsf{Pr}(A)$$

Confirmation & Probability

From Probability to Confirmation

Confirmation Measures

- Given a probability function Pr, we may construct a *confirmation measure* $\mathfrak{C}$,
  - $\mathfrak{C}(H,E)$ gives the degree to which the evidence E confirms the hypothesis H.
- One popular confirmation measure:

$$\mathfrak{D}(H,E) = \mathsf{Pr}(H \mid E) - \mathsf{Pr}(H)$$

  - There are other possibilities—*e.g.*,

$$\mathfrak{R}(H,E) = \log\left(\frac{\mathsf{Pr}(H \mid E)}{\mathsf{Pr}(H)}\right)$$

$$\mathfrak{L}(H,E) = \log\left(\frac{\mathsf{Pr}(E \mid H)}{\mathsf{Pr}(E \mid \neg H)}\right)$$

Confirmation Measures

- Given a probability function Pr, we may construct a *confirmation measure* $\mathfrak{C}$,
    - $\mathfrak{C}(H, E)$ gives the degree to which the evidence E confirms the hypothesis H.
- One popular confirmation measure:

$$\mathfrak{D}(H, E) = \mathsf{Pr}(H \mid E) - \mathsf{Pr}(H)$$

    - There are other possibilities—*e.g.*,

$$\mathfrak{R}(H, E) = \log\left(\frac{\mathsf{Pr}(H \mid E)}{\mathsf{Pr}(H)}\right)$$

$$\mathfrak{L}(H, E) = \log\left(\frac{\mathsf{Pr}(E \mid H)}{\mathsf{Pr}(E \mid \neg H)}\right)$$

Confirmation Measures

- Given a probability function Pr, we may construct a *confirmation measure* $\mathfrak{C}$,
  - $\mathfrak{C}(H, E)$ gives the degree to which the evidence E confirms the hypothesis H.
- One popular confirmation measure:

$$\mathfrak{D}(H, E) = \mathsf{Pr}(H \mid E) - \mathsf{Pr}(H)$$

  - There are other possibilities—*e.g.*,

$$\mathfrak{R}(H, E) = \log\left(\frac{\mathsf{Pr}(H \mid E)}{\mathsf{Pr}(H)}\right)$$

$$\mathfrak{L}(H, E) = \log\left(\frac{\mathsf{Pr}(E \mid H)}{\mathsf{Pr}(E \mid \neg H)}\right)$$

Confirmation Measures

- Given a probability function Pr, we may construct a *confirmation measure* $\mathfrak{C}$,
    - $\mathfrak{C}(H, E)$ gives the degree to which the evidence E confirms the hypothesis H.
- One popular confirmation measure:

$$\mathfrak{D}(H, E) = \mathsf{Pr}(H \mid E) - \mathsf{Pr}(H)$$

- There are other possibilities—*e.g.*,

$$\mathfrak{R}(H, E) = \log\left(\frac{\mathsf{Pr}(H \mid E)}{\mathsf{Pr}(H)}\right)$$

$$\mathfrak{L}(H, E) = \log\left(\frac{\mathsf{Pr}(E \mid H)}{\mathsf{Pr}(E \mid \neg H)}\right)$$

Confirmation Measures

- All of these measures will agree about the following:
    - If $\Pr(H \mid E) > \Pr(H)$, then E confirms H
    - If $\Pr(H \mid E) < \Pr(H)$, then E disconfirms H
    - If $\Pr(H \mid E) = \Pr(H)$, then E neither confirms nor disconfirms H

Confirmation Measures

- All of these measures will agree about the following:
    - If $\mathsf{Pr}(H \mid E) > \mathsf{Pr}(H)$, then E confirms H
    - If $\mathsf{Pr}(H \mid E) < \mathsf{Pr}(H)$, then E disconfirms H
    - If $\mathsf{Pr}(H \mid E) = \mathsf{Pr}(H)$, then E neither confirms nor disconfirms H

Confirmation Measures

- All of these measures will agree about the following:
  - If $\mathsf{Pr}(H \mid E) > \mathsf{Pr}(H)$, then E confirms H
  - If $\mathsf{Pr}(H \mid E) < \mathsf{Pr}(H)$, then E disconfirms H
  - If $\mathsf{Pr}(H \mid E) = \mathsf{Pr}(H)$, then E neither confirms nor disconfirms H

Confirmation Measures

- All of these measures will agree about the following:
    - If $\mathsf{Pr}(H \mid E) > \mathsf{Pr}(H)$, then E confirms H
    - If $\mathsf{Pr}(H \mid E) < \mathsf{Pr}(H)$, then E disconfirms H
    - If $\mathsf{Pr}(H \mid E) = \mathsf{Pr}(H)$, then E neither confirms nor disconfirms H

Probability & Confirmation

- It is a consequence of the definition of conditional probability that:

$$\mathsf{Pr}(H \mid E) = \frac{\mathsf{Pr}(E \mid H)}{\mathsf{Pr}(E)} \cdot \mathsf{Pr}(H)$$

- So, we may say: E confirms H if and only if

$$\mathsf{Pr}(H \mid E) > \mathsf{Pr}(H)$$

- That is: E confirms H if and only if H did a good job *predicting E*.

Probability & Confirmation

- It is a consequence of the definition of conditional probability that:

$$\mathsf{Pr}(H \mid E) = \frac{\mathsf{Pr}(E \mid H)}{\mathsf{Pr}(E)} \cdot \mathsf{Pr}(H)$$

- So, we may say: *E* confirms *H* if and only if

$$\mathsf{Pr}(H \mid E) > \mathsf{Pr}(H)$$

- That is: *E* confirms *H* if and only if *H* did a good job *predicting E*.

Probability & Confirmation

- It is a consequence of the definition of conditional probability that:

$$\Pr(H \mid E) = \frac{\Pr(E \mid H)}{\Pr(E)} \cdot \Pr(H)$$

- So, we may say: E confirms H if and only if

$$\frac{\Pr(E \mid H)}{\Pr(E)} \cdot \Pr(H) > \Pr(H)$$

- That is: E confirms H if and only if H did a good job *predicting E*.

Probability & Confirmation

- It is a consequence of the definition of conditional probability that:

$$\mathsf{Pr}(H \mid E) = \frac{\mathsf{Pr}(E \mid H)}{\mathsf{Pr}(E)} \cdot \mathsf{Pr}(H)$$

- So, we may say: *E* confirms *H* if and only if

$$\mathsf{Pr}(E \mid H) > \mathsf{Pr}(E)$$

- That is: *E* confirms *H* if and only if *H* did a good job *predicting E*.

Probability & Confirmation

- It is a consequence of the definition of conditional probability that:

$$\mathsf{Pr}(H \mid E) = \frac{\mathsf{Pr}(E \mid H)}{\mathsf{Pr}(E)} \cdot \mathsf{Pr}(H)$$

- So, we may say: *E* confirms *H* if and only if

$$\mathsf{Pr}(E \mid H) > \mathsf{Pr}(E)$$

- That is: *E* confirms *H* if and only if *H* did a good job *predicting E*.

Probability & Confirmation

- E confirms H if and only if H did a good job predicting E.
- What do we mean by 'good job'?
  - In order to do a good job predicting E, H doesn't have to make E likely.
  - Also, in order to do a good job predicting E, it is not enough for H to make E likely.
  - To do a good job predicting E, H must make E more likely than its negation, $\neg H$.

Probability & Confirmation

- E confirms H if and only if H did a good job predicting E.
- What do we mean by 'good job'?
    - In order to do a good job predicting E, H doesn't have to make E likely.
    - Also, in order to do a good job predicting E, it is not enough for H to make E likely.
    - To do a good job predicting E, H must make E more likely than its negation, $\neg H$.

Probability & Confirmation

- E confirms H if and only if H did a good job predicting E.
- What do we mean by 'good job'?
    - In order to do a good job predicting E, H doesn't have to make E likely.
    - Also, in order to do a good job predicting E, it is not enough for H to make E likely.
    - To do a good job predicting E, H must make E more likely than its negation, $\neg H$.

Probability & Confirmation

- E confirms H if and only if H did a good job predicting E.
- What do we mean by 'good job'?
    - In order to do a good job predicting E, H doesn't have to make E likely.
    - Also, in order to do a good job predicting E, it is not enough for H to make E likely.
    - To do a good job predicting E, H must make E more likely than its negation, $\neg H$.

Probability & Confirmation

- E confirms H if and only if H did a good job predicting E.
- What do we mean by 'good job'?
    - In order to do a good job predicting E, H doesn't have to make E likely.
    - Also, in order to do a good job predicting E, it is not enough for H to make E likely.
    - To do a good job predicting E, H must make E more likely than its negation, $\neg H$.

Probability & Confirmation

- The suspect confesses.
    - This is *very* unlikely, given that the suspect is guilty. Guilty suspects almost never confess.
    - However, it is even *less* likely that the suspect confesses, given that the suspect is innocent.
    - So, a confession confirms the hypothesis that the suspect was guilty, even though a confession was *very* unlikely given that the suspect was guilty.

Probability & Confirmation

- The suspect confesses.
  - This is *very* unlikely, given that the suspect is guilty. Guilty suspects almost never confess.
  - However, it is even *less* likely that the suspect confesses, given that the suspect is innocent.
  - So, a confession confirms the hypothesis that the suspect was guilty, even though a confession was *very* unlikely given that the suspect was guilty.

Probability & Confirmation

- The suspect confesses.
  - This is *very* unlikely, given that the suspect is guilty. Guilty suspects almost never confess.
  - However, it is even *less* likely that the suspect confesses, given that the suspect is innocent.
  - So, a confession confirms the hypothesis that the suspect was guilty, even though a confession was *very* unlikely given that the suspect was guilty.

Probability & Confirmation

- The suspect confesses.
    - This is *very* unlikely, given that the suspect is guilty. Guilty suspects almost never confess.
    - However, it is even *less* likely that the suspect confesses, given that the suspect is innocent.
    - So, a confession confirms the hypothesis that the suspect was guilty, even though a confession was *very* unlikely given that the suspect was guilty.

Probability & Confirmation

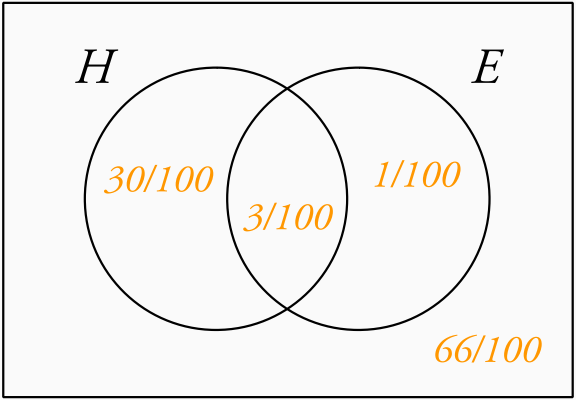

Probability & Confirmation

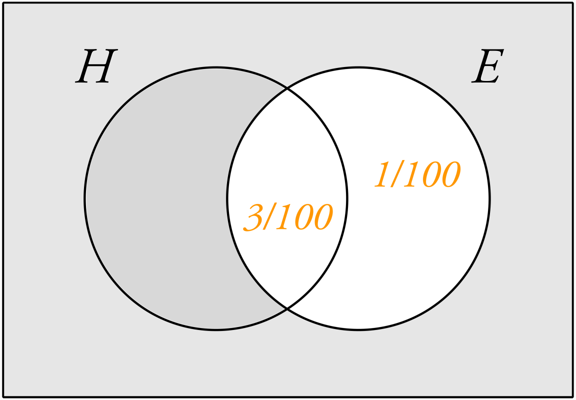

Probability & Confirmation

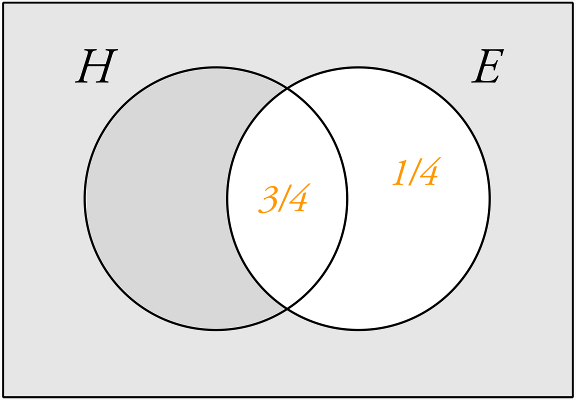

Probability & Confirmation

- Hypothesis, H: Bob wears a helmet on his bike ride into work.
- Evidence, E: Bob makes it into work without getting into an accident.
- Given H, E is *very* likely.
- However, E is even *more* likely given $\neg H$.
- So E doesn't confirm H, even though H made E very likely.

Probability & Confirmation

- Hypothesis, H: Bob wears a helmet on his bike ride into work.
- Evidence, E: Bob makes it into work without getting into an accident.
- Given H, E is *very* likely.
- However, E is even *more* likely given $\neg H$.
- So E doesn't confirm H, even though H made E very likely.

Probability & Confirmation

- Hypothesis, H: Bob wears a helmet on his bike ride into work.
- Evidence, E: Bob makes it into work without getting into an accident.
- Given H, E is *very* likely.
- However, E is even *more* likely given $\neg H$.
- So E doesn't confirm H, even though H made E very likely.

Probability & Confirmation

- Hypothesis, H: Bob wears a helmet on his bike ride into work.
- Evidence, E: Bob makes it into work without getting into an accident.
- Given H, E is *very* likely.
- However, E is even *more* likely given $\neg H$.
- So E doesn't confirm H, even though H made E very likely.

Probability & Confirmation

- Hypothesis, H: Bob wears a helmet on his bike ride into work.
- Evidence, E: Bob makes it into work without getting into an accident.
- Given H, E is *very* likely.
- However, E is even *more* likely given $\neg H$.
- So E doesn't confirm H, even though H made E very likely.

Probability & Confirmation

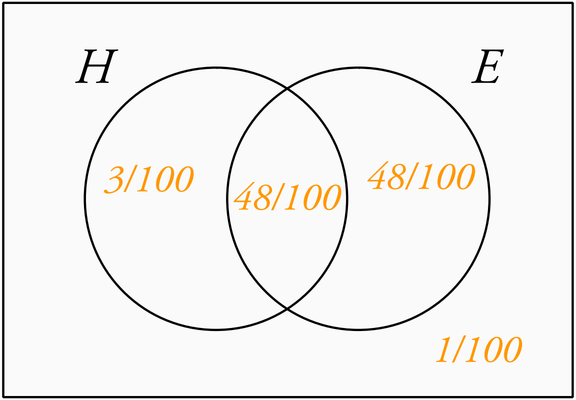

Probability & Confirmation

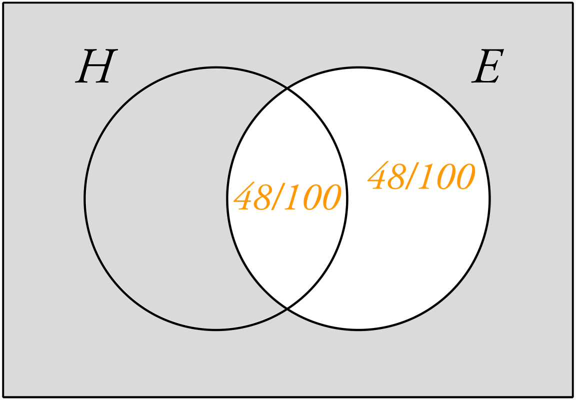

Probability & Confirmation

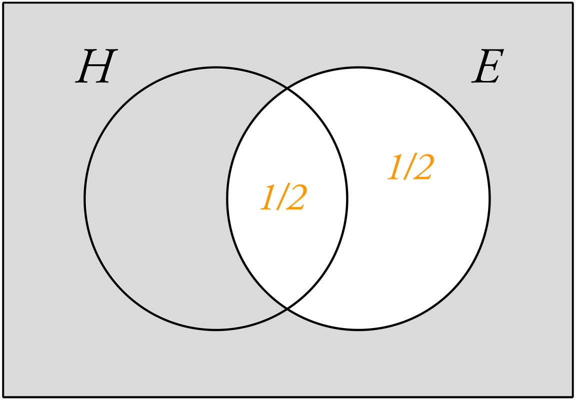

Break

Bayesian Confirmation Theory

Credences

- Our theory of confirmation says nothing until we say more about the probability function Pr.
- The Bayesian thinks that Pr represents some hypothetical rational agent's *degrees of belief*, or *credences*.
    - If $\Pr(A) = 1$, then the agent thinks that A is certainly true.
    - If $\Pr(A) = 0$, then the agent thinks that A is certainly false.
    - If $\Pr(A) = 1/2$, then the agent is as confident that A is true as they are that A is false.

Credences

- Our theory of confirmation says nothing until we say more about the probability function Pr.
- The Bayesian thinks that Pr represents some hypothetical rational agent's *degrees of belief*, or *credences*.
    - If $\mathsf{Pr}(A) = 1$, then the agent thinks that A is certainly true.
    - If $\mathsf{Pr}(A) = 0$, then the agent thinks that A is certainly false.
    - If $\mathsf{Pr}(A) = 1/2$, then the agent is as confident that A is true as they are that A is false.

Credences

- Our theory of confirmation says nothing until we say more about the probability function Pr.
- The Bayesian thinks that Pr represents some hypothetical rational agent's *degrees of belief*, or *credences*.
    - If $\mathsf{Pr}(A) = 1$, then the agent thinks that A is certainly true.
    - If $\mathsf{Pr}(A) = 0$, then the agent thinks that A is certainly false.
    - If $\mathsf{Pr}(A) = 1/2$, then the agent is as confident that A is true as they are that A is false.

Credences

- Our theory of confirmation says nothing until we say more about the probability function Pr.
- The Bayesian thinks that Pr represents some hypothetical rational agent's *degrees of belief*, or *credences*.
    - If $\mathsf{Pr}(A) = 1$, then the agent thinks that A is certainly true.
    - If $\mathsf{Pr}(A) = 0$, then the agent thinks that A is certainly false.
    - If $\mathsf{Pr}(A) = 1/2$, then the agent is as confident that A is true as they are that A is false.

Credences

- Our theory of confirmation says nothing until we say more about the probability function Pr.
- The Bayesian thinks that Pr represents some hypothetical rational agent's *degrees of belief*, or *credences*.
    - If $\mathsf{Pr}(A) = 1$, then the agent thinks that A is certainly true.
    - If $\mathsf{Pr}(A) = 0$, then the agent thinks that A is certainly false.
    - If $\mathsf{Pr}(A) = 1/2$, then the agent is as confident that A is true as they are that A is false.

Probabilism

- The first claim of Bayesianism:

Probabilism

It is a requirement of rationality that your degrees of belief $\Pr$ satisfy the axioms of probability.

Probabilism

- The first claim of Bayesianism:

Probabilism

It is a requirement of rationality that your degrees of belief $\Pr$ satisfy the axioms of probability.

Conditionalization

- Suppose that, upon acquiring the total evidence E, you are disposed to adopt some new degrees of belief, Pr_E.
- The second claim of Bayesianism:

Conditionalization

It is a requirement of rationality that, upon acquiring the total evidence E, you are disposed to adopt a new credence function Pr_E which is your old credence function *conditionalized on E*. That is, for all H,

$$\mathsf{Pr}_E(H) = \mathsf{Pr}(H \mid E) = \frac{\mathsf{Pr}(E \mid H)}{\mathsf{Pr}(E)} \cdot \mathsf{Pr}(H)$$

Conditionalization

- Suppose that, upon acquiring the total evidence E, you are disposed to adopt some new degrees of belief, Pr_E.
- The second claim of Bayesianism:

Conditionalization

It is a requirement of rationality that, upon acquiring the total evidence E, you are disposed to adopt a new credence function Pr_E which is your old credence function *conditionalized on E*. That is, for all H,

$$\mathsf{Pr}_E(H) = \mathsf{Pr}(H \mid E) = \frac{\mathsf{Pr}(E \mid H)}{\mathsf{Pr}(E)} \cdot \mathsf{Pr}(H)$$

Conditionalization

- Suppose that, upon acquiring the total evidence E, you are disposed to adopt some new degrees of belief, Pr_E.
- The second claim of Bayesianism:

Conditionalization

It is a requirement of rationality that, upon acquiring the total evidence E, you are disposed to adopt a new credence function Pr_E which is your old credence function *conditionalized on E*. That is, for all H,

$$\mathsf{Pr}_E(H) = \mathsf{Pr}(H \mid E) = \frac{\mathsf{Pr}(E \mid H)}{\mathsf{Pr}(E)} \cdot \mathsf{Pr}(H)$$

Priors and Posteriors

- Pr is the *prior* credence function.
- Pr_E is the *posterior* credence function.

Priors and Posteriors

- Pr is the *prior* credence function.
- Pr_E is the *posterior* credence function.

Bayesian Theory of Confirmation

- The Bayesian theory of confirmation says that E confirms H iff

$$\mathsf{Pr}_E(H) > \mathsf{Pr}(H)$$

- And E disconfirms H iff

$$\mathsf{Pr}_E(H) < \mathsf{Pr}(H)$$

Bayesian Theory of Confirmation

- The Bayesian theory of confirmation says that E confirms H iff

$$\mathsf{Pr}_E(H) > \mathsf{Pr}(H)$$

- And E disconfirms H iff

$$\mathsf{Pr}_E(H) < \mathsf{Pr}(H)$$

Justifying Bayesianism Pragmatically

- A pragmatic justification of probabilism: If your degrees of belief don't satisfy the axioms of probability, then you could be sold a combination of bets which is guaranteed to lose you money *come what may*.
- A pragmatic justification of conditionalization: If you stand to learn whether E, and you are disposed to revise your beliefs in any way other than conditionalization, then you could be reliably sold a series of bets which are guaranteed to lose you money no matter what.

Justifying Bayesianism Pragmatically

- A pragmatic justification of probabilism: If your degrees of belief don't satisfy the axioms of probability, then you could be sold a combination of bets which is guaranteed to lose you money *come what may*.
- A pragmatic justification of conditionalization: If you stand to learn whether E, and you are disposed to revise your beliefs in any way other than conditionalization, then you could be reliably sold a series of bets which are guaranteed to lose you money no matter what.

Two Questions About Induction

- Question 1: What are the canons of inductive logic?
- Question 2: Why should we think that good inductive inferences will lead us to truth?

Two Questions About Induction

- Question 1: What are the canons of inductive logic?
- Question 2: Why should we think that good inductive inferences will lead us to truth?

Justifying Bayesianism Alethically

- An alethic justification of probabilism: If your degrees of belief don't satisfy the axioms of probability, then there is some other degrees of belief you could adopt which is guaranteed to be *more accurate* than yours, no matter what.
- An alethic justification of conditionalization: If you stand to learn whether E, then the strategy of conditionalization has higher *expected accuracy* than any other strategy of belief-revision.

Justifying Bayesianism Alethically

- An alethic justification of probabilism: If your degrees of belief don't satisfy the axioms of probability, then there is some other degrees of belief you could adopt which is guaranteed to be *more accurate* than yours, no matter what.
- An alethic justification of conditionalization: If you stand to learn whether E, then the strategy of conditionalization has higher *expected accuracy* than any other strategy of belief-revision.

Why the Bayesian Thinks You Can't Always Get What You Want

Hempel's Impossibility Results, Again

Entailments Confirm

If *H entails E*, then *E confirms H*

Consequence Condition

If *E* confirms *H*, then *E* confirms anything which *H* entails.

- If we accept both Entailments Confirm and Consequence Condition, then every proposition confirms every other proposition.

Hempel's Impossibility Results, Again

Entailments Confirm

If *H entails E*, then *E confirms H*

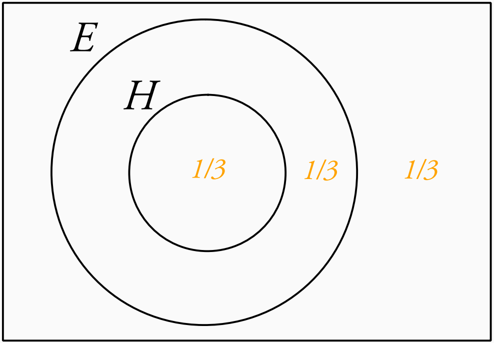

Hempel's Impossibility Results, Again

Entailments Confirm

If H *entails* E, then E *confirms* H

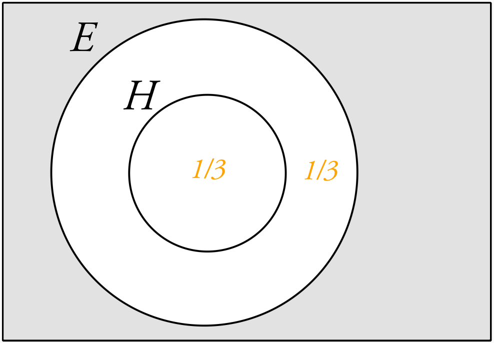

Hempel's Impossibility Results, Again

Entailments Confirm

If H *entails* E, then E *confirms* H

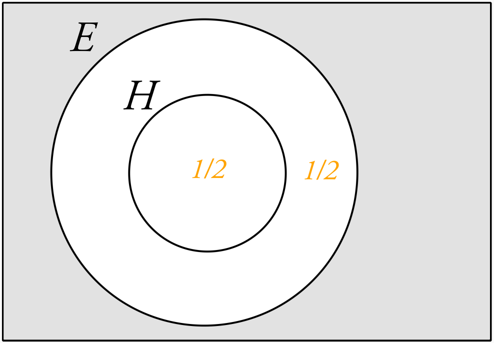

Hempel's Impossibility Results, Again

Consequence Condition

If E confirms H, then E confirms anything which H entails.

- We are playing poker, and I catch a glimpse of your cards. I see that you have a spade.
- That you have a spade confirms that you have the ace of spades.
- That you have the ace of spades entails that you have an ace.
- But that you have a spade does not confirm that you have an ace.

Hempel's Impossibility Results, Again

Consequence Condition

If E confirms H, then E confirms anything which H entails.

- We are playing poker, and I catch a glimpse of your cards. I see that you have a spade.
- That you have a spade confirms that you have the ace of spades.
- That you have the ace of spades entails that you have an ace.
- But that you have a spade does not confirm that you have an ace.

Hempel's Impossibility Results, Again

Consequence Condition

If E confirms H, then E confirms anything which H entails.

- We are playing poker, and I catch a glimpse of your cards. I see that you have a spade.
- That you have a spade confirms that you have the ace of spades.
- That you have the ace of spades entails that you have an ace.
- But that you have a spade does not confirm that you have an ace.

Hempel's Impossibility Results, Again

Consequence Condition

If E confirms H, then E confirms anything which H entails.

- We are playing poker, and I catch a glimpse of your cards. I see that you have a spade.
- That you have a spade confirms that you have the ace of spades.
- That you have the ace of spades entails that you have an ace.
- But that you have a spade does not confirm that you have an ace.

Hempel's Impossibility Results, Again

Consequence Condition

If E confirms H, then E confirms anything which H entails.

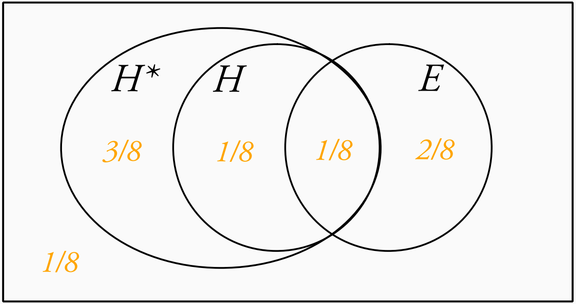

Hempel's Impossibility Results, Again

Consequence Condition

If E confirms H, then E confirms anything which H entails.

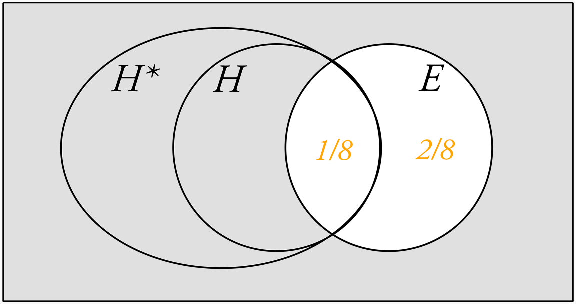

Hempel's Impossibility Results, Again

Consequence Condition

If E confirms H, then E confirms anything which H entails.

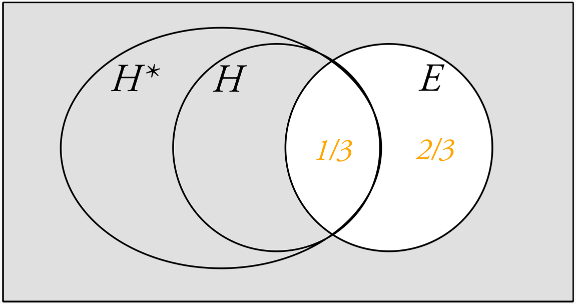

Hempel's Impossibility Results, Again

Laws are Confirmed by Their Instances

A law statement of the form "All *F*s are *G*s" is confirmed by an *F* *G*

Equivalence Condition

If E confirms H, then E confirms anything equivalent to H

- If we accept both of these principles, then we must say that a green leaf confirms the hypothesis that "All ravens are black".

Hempel's Impossibility Results, Again

Equivalence Condition

If E confirms H, then E confirms anything equivalent to H

- If H is equivalent to H^*, then $\Pr(H) = \Pr(H^*)$ and $\Pr(H \& E) = \Pr(H^* \& E)$. So

$$\Pr(H^* \mid E) = \frac{\Pr(H^* \& E)}{\Pr(E)} = \frac{\Pr(H \& E)}{\Pr(E)} = \Pr(H \mid E)$$

Hempel's Impossibility Results, Again

Equivalence Condition

If E confirms H, then E confirms anything equivalent to H

- If H is equivalent to H^*, then $\Pr(H) = \Pr(H^*)$ and $\Pr(H \& E) = \Pr(H^* \& E)$. So

$$\Pr(H^* \mid E) = \frac{\Pr(H^* \& E)}{\Pr(E)} = \frac{\Pr(H \& E)}{\Pr(E)} = \Pr(H \mid E)$$

Hempel's Impossibility Results, Again

Laws are Confirmed by Their Instances

A law of the form "All *F*s are *G*s" is confirmed by an *F G*

<u>All</u>

$$\begin{array}{r}\\ \text{Raven} \\ \text{Non-Raven}\end{array}\begin{array}{c}\begin{array}{cc}\text{Black} & \text{Non-Black}\end{array}\\ \begin{bmatrix} 4 & 0 \\ 2 & 2 \end{bmatrix}\end{array}$$

<u>Some</u>

$$\begin{array}{r}\\ \text{Raven} \\ \text{Non-Raven}\end{array}\begin{array}{c}\begin{array}{cc}\text{Black} & \text{Non-Black}\end{array}\\ \begin{bmatrix} 2 & 2 \\ 2 & 2 \end{bmatrix}\end{array}$$

Hempel's Impossibility Results, Again

- $E =$ a randomly selected thing is a non-black non-raven.
- As we saw, E will confirm *All* iff *All* makes E more likely than *Some*. But

$$\Pr(E \mid All) = 1/4 \qquad \text{and} \qquad \Pr(E \mid Some) = 1/4$$

- So the Universal hypothesis *All* is not confirmed by a non-black non-raven.
- So Laws are Confirmed by Their Instances is false.

Hempel's Impossibility Results, Again

- $E =$ a randomly selected thing is a non-black non-raven.
- As we saw, E will confirm *All* iff *All* makes E more likely than *Some*. But

$$\Pr(E \mid All) = 1/4 \qquad \text{and} \qquad \Pr(E \mid Some) = 1/4$$

- So the Universal hypothesis *All* is not confirmed by a non-black non-raven.
- So Laws are Confirmed by Their Instances is false.

Hempel's Impossibility Results, Again

- $E =$ a randomly selected thing is a non-black non-raven.
- As we saw, E will confirm *All* iff *All* makes E more likely than *Some*. But

$$\Pr(E \mid All) = 1/4 \qquad \text{and} \qquad \Pr(E \mid Some) = 1/4$$

- So the Universal hypothesis *All* is not confirmed by a non-black non-raven.
- So Laws are Confirmed by Their Instances is false.

Hempel's Impossibility Results, Again

- E = a randomly selected thing is a non-black non-raven.
- As we saw, E will confirm *All* iff *All* makes E more likely than *Some*. But

$$\mathsf{Pr}(E \mid All) = 1/4 \qquad \text{and} \qquad \mathsf{Pr}(E \mid Some) = 1/4$$

- So the Universal hypothesis *All* is not confirmed by a non-black non-raven.
- So Laws are Confirmed by Their Instances is false.

Hempel's Impossibility Results, Again

$$\underline{\text{All}} \qquad\qquad\qquad\qquad \underline{\text{Some}}$$

$$\begin{array}{rcc} & \text{Black} & \text{Non-Black} \\ \text{Raven} & 4 & 0 \\ \text{Non-Raven} & 2 & 2 \end{array} \qquad \begin{array}{rcc} & \text{Black} & \text{Non-Black} \\ \text{Raven} & 2 & 2 \\ \text{Non-Raven} & 2 & 2 \end{array}$$

- E^* = A randomly selected thing is a black raven.
- As we saw, E^* will confirm *All* iff *All* makes E^* more likely than *Some* does. And

$$\Pr(E^* \mid All) = 1/2 \qquad \text{and} \qquad \Pr(E^* \mid Some) = 1/4$$

- So a black raven confirms *All*, even though a non-black non-raven does not.

Hempel's Impossibility Results, Again

<u>All</u>

$$\begin{array}{c|cc} & \text{Black} & \text{Non-Black} \\ \hline \text{Raven} & 4 & 0 \\ \text{Non-Raven} & 2 & 2 \end{array}$$

<u>Some</u>

$$\begin{array}{c|cc} & \text{Black} & \text{Non-Black} \\ \hline \text{Raven} & 2 & 2 \\ \text{Non-Raven} & 2 & 2 \end{array}$$

- $E^* =$ A randomly selected thing is a black raven.
- As we saw, E^* will confirm *All* iff *All* makes E^* more likely than *Some* does. And

$$\mathsf{Pr}(E^* \mid All) = 1/2 \qquad \text{and} \qquad \mathsf{Pr}(E^* \mid Some) = 1/4$$

- So a black raven confirms *All*, even though a non-black non-raven does not.

Hempel's Impossibility Results, Again

$$\begin{array}{c} \underline{\text{All}} \\ \begin{array}{r} \\ \text{Raven} \\ \text{Non-Raven} \end{array} \begin{array}{c} \begin{array}{cc} \text{Black} & \text{Non-Black} \end{array} \\ \begin{bmatrix} 4 & 0 \\ 2 & 2 \end{bmatrix} \end{array} \end{array} \qquad \begin{array}{c} \underline{\text{Some}} \\ \begin{array}{r} \\ \text{Raven} \\ \text{Non-Raven} \end{array} \begin{array}{c} \begin{array}{cc} \text{Black} & \text{Non-Black} \end{array} \\ \begin{bmatrix} 2 & 2 \\ 2 & 2 \end{bmatrix} \end{array} \end{array}$$

- E^* = A randomly selected thing is a black raven.
- As we saw, E^* will confirm *All* iff *All* makes E^* more likely than *Some* does. And

$$\Pr(E^* \mid \mathit{All}) = 1/2 \qquad \text{and} \qquad \Pr(E^* \mid \mathit{Some}) = 1/4$$

- So a black raven confirms *All*, even though a non-black non-raven does not.

Goodman's Impossibility Result

- *Green* = All emeralds are green
- *Grue* = All emeralds are grue
- *E* = All observed emeralds are green/grue

Goodman's Impossibility Result

- *Green* = All emeralds are green
- *Grue* = All emeralds are grue
- *E* = All observed emeralds are green/grue

Goodman's Impossibility Result

- *Green* = All emeralds are green
- *Grue* = All emeralds are grue
- *E* = All observed emeralds are green/grue

Goodman's Impossibility Result

$$\frac{\mathsf{Pr}(Green \mid E)}{\mathsf{Pr}(Grue \mid E)} = \frac{\frac{\mathsf{Pr}(E|Green)}{\mathsf{Pr}(E)} \cdot \mathsf{Pr}(Green)}{\frac{\mathsf{Pr}(E|Grue)}{\mathsf{Pr}(E)} \cdot \mathsf{Pr}(Grue)}$$

$$= \frac{\mathsf{Pr}(E \mid Green) \cdot \mathsf{Pr}(Green)}{\mathsf{Pr}(E \mid Grue) \cdot \mathsf{Pr}(Grue)}$$

$$= \frac{\mathsf{Pr}(Green)}{\mathsf{Pr}(Grue)}$$

The Problem of the Priors

The Problem of the Priors

- If we want *Green* to have a higher posterior credence than *Grue*, then we must stipulate that *Green* has a higher *prior* credence than *Grue*.
- The "problem of the priors": which prior credences is it rational to adopt?

The Problem of the Priors

- If we want *Green* to have a higher posterior credence than *Grue*, then we must stipulate that *Green* has a higher *prior* credence than *Grue*.
- The "problem of the priors": which prior credences is it rational to adopt?

What We Want from a Theory of Confirmation

- A *qualitative* account of confirmation.
    - For any H, E: does E confirm H?
- A *quantitative* measure of confirmation.
    - For any H, E: to what *degree* does E confirm H?
- We'd like our theory of confirmation to be ~~*formal* and~~ *intersubjective*.
    - ~~*Formal*: we can say whether E confirms H by looking only at syntax, or logical *form*.~~
    - *Intersubjective*: we can all agree about whether E confirms H

Subjectivism *vs* Objectivism

- Radical Subjectivism: All probabilistic priors are rationally permissible.
- Only slightly less radical subjectivism: Any probabilistic prior is rationally permissible so long as it satisfies a PROBABILITY COORDINATION PRINCIPLE like

$$\text{if } H \text{ gives } E \text{ an objective chance of } x \text{, then } \Pr(E \mid H) = x$$

- Moderate Subjectivism: There is a limited range of rationally permissible priors.
- Objectivism: There is only one rational prior.

Subjectivism *vs* Objectivism

- Radical Subjectivism: All probabilistic priors are rationally permissible.
- Only slightly less radical subjectivism: Any probabilistic prior is rationally permissible so long as it satisfies a PROBABILITY COORDINATION PRINCIPLE like

$$\text{if } H \text{ gives } E \text{ an objective chance of } x\text{, then } \Pr(E \mid H) = x$$

- Moderate Subjectivism: There is a limited range of rationally permissible priors.
- Objectivism: There is only one rational prior.

Subjectivism *vs* Objectivism

- Radical Subjectivism: All probabilistic priors are rationally permissible.
- Only slightly less radical subjectivism: Any probabilistic prior is rationally permissible so long as it satisfies a PROBABILITY COORDINATION PRINCIPLE like

$$\text{if } H \text{ gives } E \text{ an objective chance of } x\text{, then } \mathsf{Pr}(E \mid H) = x$$

- Moderate Subjectivism: There is a limited range of rationally permissible priors.
- Objectivism: There is only one rational prior.

Subjectivism *vs* Objectivism

- Radical Subjectivism: All probabilistic priors are rationally permissible.
- Only slightly less radical subjectivism: Any probabilistic prior is rationally permissible so long as it satisfies a PROBABILITY COORDINATION PRINCIPLE like

$$\text{if } H \text{ gives } E \text{ an objective chance of } x\text{, then } \mathsf{Pr}(E \mid H) = x$$

- Moderate Subjectivism: There is a limited range of rationally permissible priors.
- Objectivism: There is only one rational prior.

The Principle of Indifference

- The *Principle of Indifference* undergirds one version of Objectivism.

The Principle of Indifference

In the absence of evidence, assume a *uniform* credence distribution.

The Principle of Indifference

- The *Principle of Indifference* undergirds one version of Objectivism.

The Principle of Indifference

In the absence of evidence, assume a *uniform* credence distribution.

Bertrand's Paradox

- I drove 2100 miles from Pittsburgh to L.A. The trip took somewhere between 30 and 42 hours.
- What is the rational credence to have that it took between 30 and 35 hours?
- Principle of Indifference:

$$\Pr(30 \leq t \leq 35) = \frac{5}{12}$$

Bertrand's Paradox

- I drove 2100 miles from Pittsburgh to L.A. The trip took somewhere between 30 and 42 hours.
- What is the rational credence to have that it took between 30 and 35 hours?
- Principle of Indifference:

$$\Pr(30 \leq t \leq 35) = \frac{5}{12}$$

Bertrand's Paradox

- I drove 2100 miles from Pittsburgh to L.A. The trip took somewhere between 30 and 42 hours.
- What is the rational credence to have that it took between 30 and 35 hours?
- Principle of Indifference:

$$\Pr(30 \leq t \leq 35) = \frac{5}{12}$$

Bertrand's Paradox

- I drove 2100 miles from Pittsburgh to L.A. My average velocity was somewhere between 50 and 70 mph.
- What is the rational credence to have that the average velocity was between 60 and 70 mph?
- Principle of Indifference:

$$\Pr(60 \leq v \leq 70) = \frac{1}{2}$$

Bertrand's Paradox

- I drove 2100 miles from Pittsburgh to L.A. My average velocity was somewhere between 50 and 70 mph.
- What is the rational credence to have that the average velocity was between 60 and 70 mph?
- Principle of Indifference:

$$\Pr(60 \leq v \leq 70) = \frac{1}{2}$$

Bertrand's Paradox

- I drove 2100 miles from Pittsburgh to L.A. My average velocity was somewhere between 50 and 70 mph.
- What is the rational credence to have that the average velocity was between 60 and 70 mph?
- Principle of Indifference:

$$\Pr(60 \leq v \leq 70) = \frac{1}{2}$$

Bertrand's Paradox

- The two cases are exactly the same.
  - Principle of Indifference: $\Pr(30 \leq t \leq 35) = \frac{5}{12}$
  - Principle of Indifference: $\Pr(60 \leq v \leq 70) = \frac{1}{2}$
- But $30 \leq t \leq 35$ if and only if $60 \leq v \leq 70$. So they must receive the same credence, assuming Probabilism.

Bertrand's Paradox

- The two cases are exactly the same.
    - Principle of Indifference: $\Pr(30 \leq t \leq 35) = \frac{5}{12}$
    - Principle of Indifference: $\Pr(60 \leq v \leq 70) = \frac{1}{2}$
- But $30 \leq t \leq 35$ if and only if $60 \leq v \leq 70$. So they must receive the same credence, assuming Probabilism.

Bertrand's Paradox

- The two cases are exactly the same.
  - Principle of Indifference: $\mathsf{Pr}(30 \leq t \leq 35) = \frac{5}{12}$
  - Principle of Indifference: $\mathsf{Pr}(60 \leq v \leq 70) = \frac{1}{2}$
- But $30 \leq t \leq 35$ if and only if $60 \leq v \leq 70$. So they must receive the same credence, assuming Probabilism.

Bertrand's Paradox

- The two cases are exactly the same.
  - Principle of Indifference: $\Pr(30 \leq t \leq 35) = \frac{5}{12}$
  - Principle of Indifference: $\Pr(60 \leq v \leq 70) = \frac{1}{2}$
- But $30 \leq t \leq 35$ if and only if $60 \leq v \leq 70$. So they must receive the same credence, assuming Probabilism.

What We Want from a Theory of Confirmation

- A *qualitative* account of confirmation.
    - For any H, E: does E confirm H?
- A *quantitative* measure of confirmation.
    - For any H, E: to what *degree* does E confirm H?
- ~~We'd like our theory of confirmation to be *formal* and *intersubjective*.~~
    - ~~*Formal*: we can say whether E confirms H by looking only at syntax, or logical *form*.~~
    - ~~*Intersubjective*: we can all agree about whether E confirms H~~

Questions?