Priors, Total Probability, Expectation, Multiple Trials

CS 2800: Discrete Structures, Fall 2014

Sid Chaudhuri

Bayes' Theorem

- Given: prior probabilities of hypotheses, and the probability that each hypothesis produces the observed evidence
- Produces: probability of a particular hypothesis, given the observed evidence.

$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

Estimating $\mathrm{P}(B)$

- **Total Probability Theorem:** If $A_1, A_2, A_3, \ldots$ is a (finite or countably infinite) partition of S (they are pairwise disjoint subsets and their union is S), and B is any event in S, then

$$
\begin{aligned}
\mathrm{P}(B) &= \mathrm{P}(\cup_i (B \cap A_i)) && \text{(Set theory)} \\
&= \textstyle\sum_i \mathrm{P}(B \cap A_i) && \text{(Axiom 3)} \\
&= \textstyle\sum_i \mathrm{P}(B \mid A_i)\, \mathrm{P}(A_i) && \text{(Definition)}
\end{aligned}
$$

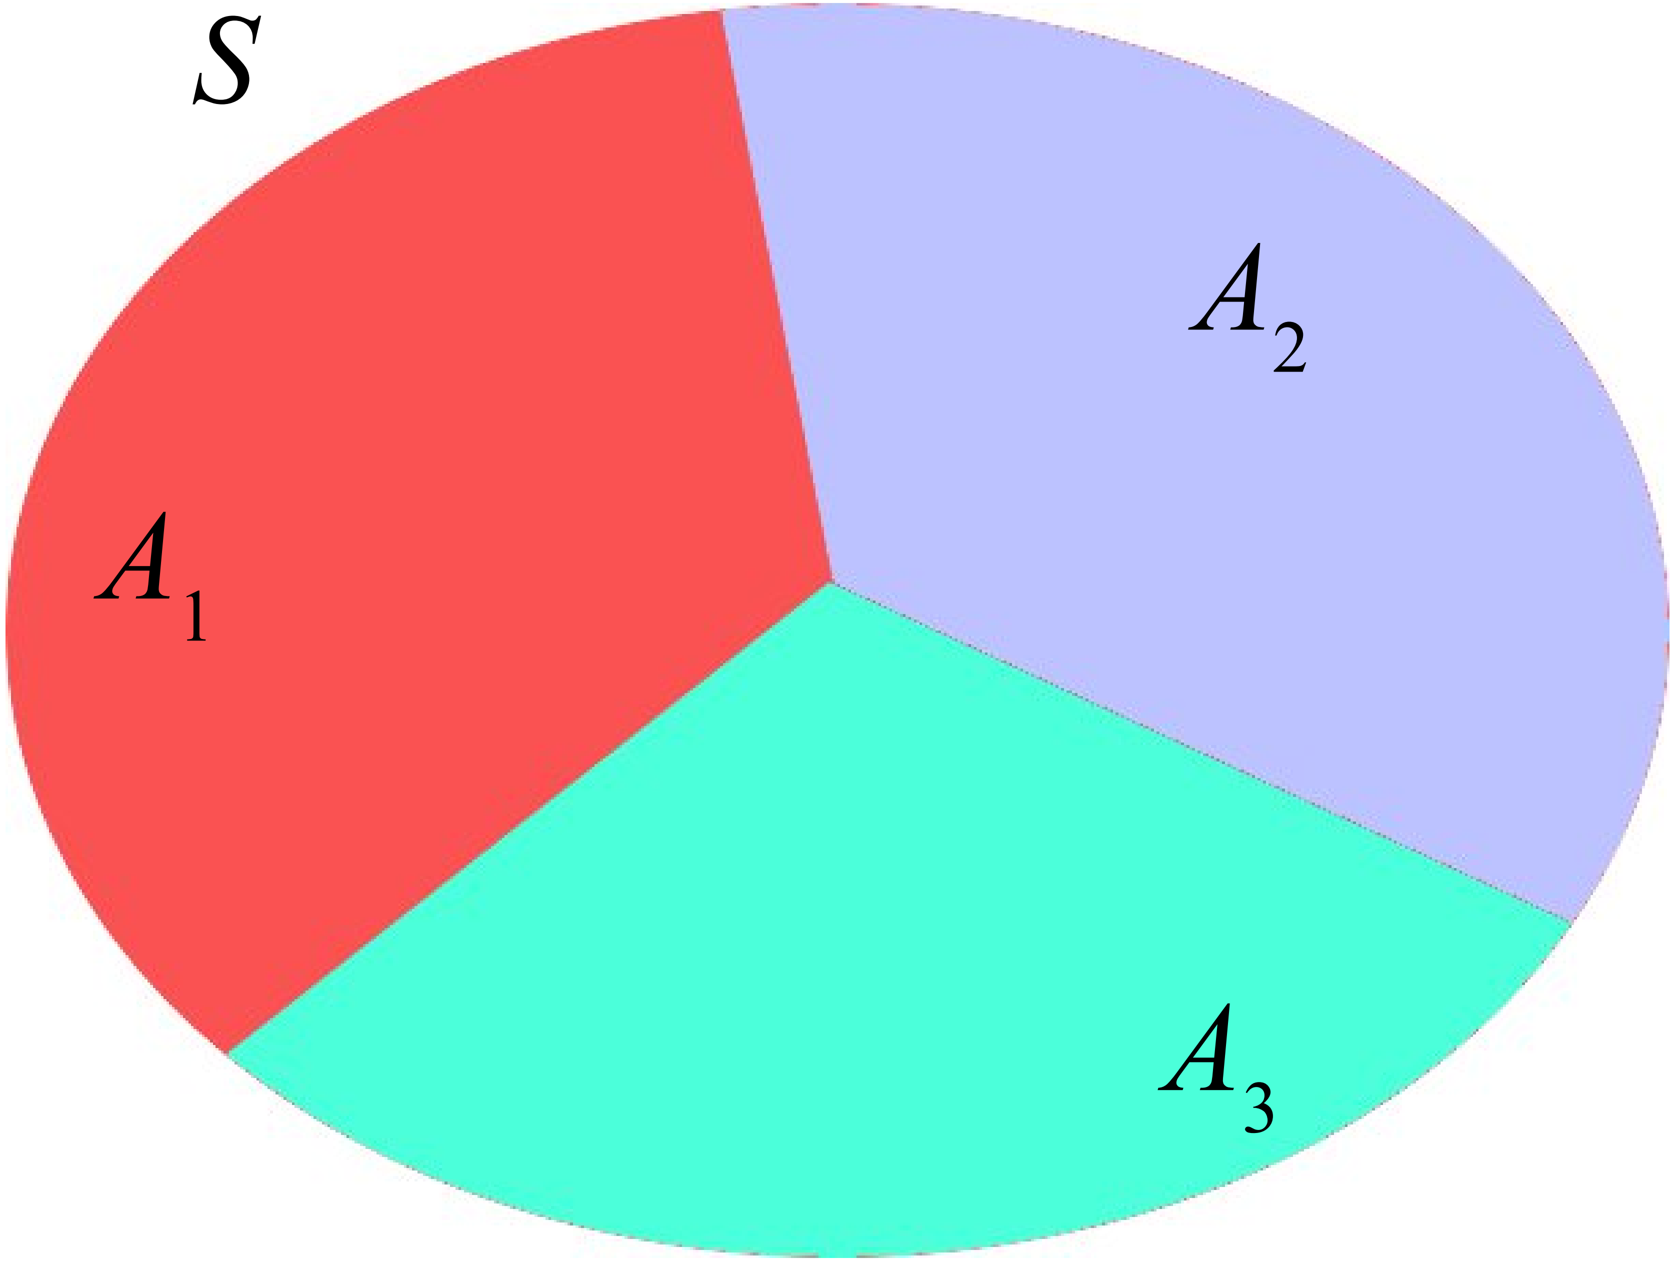

S
A_2
A_1
A_3

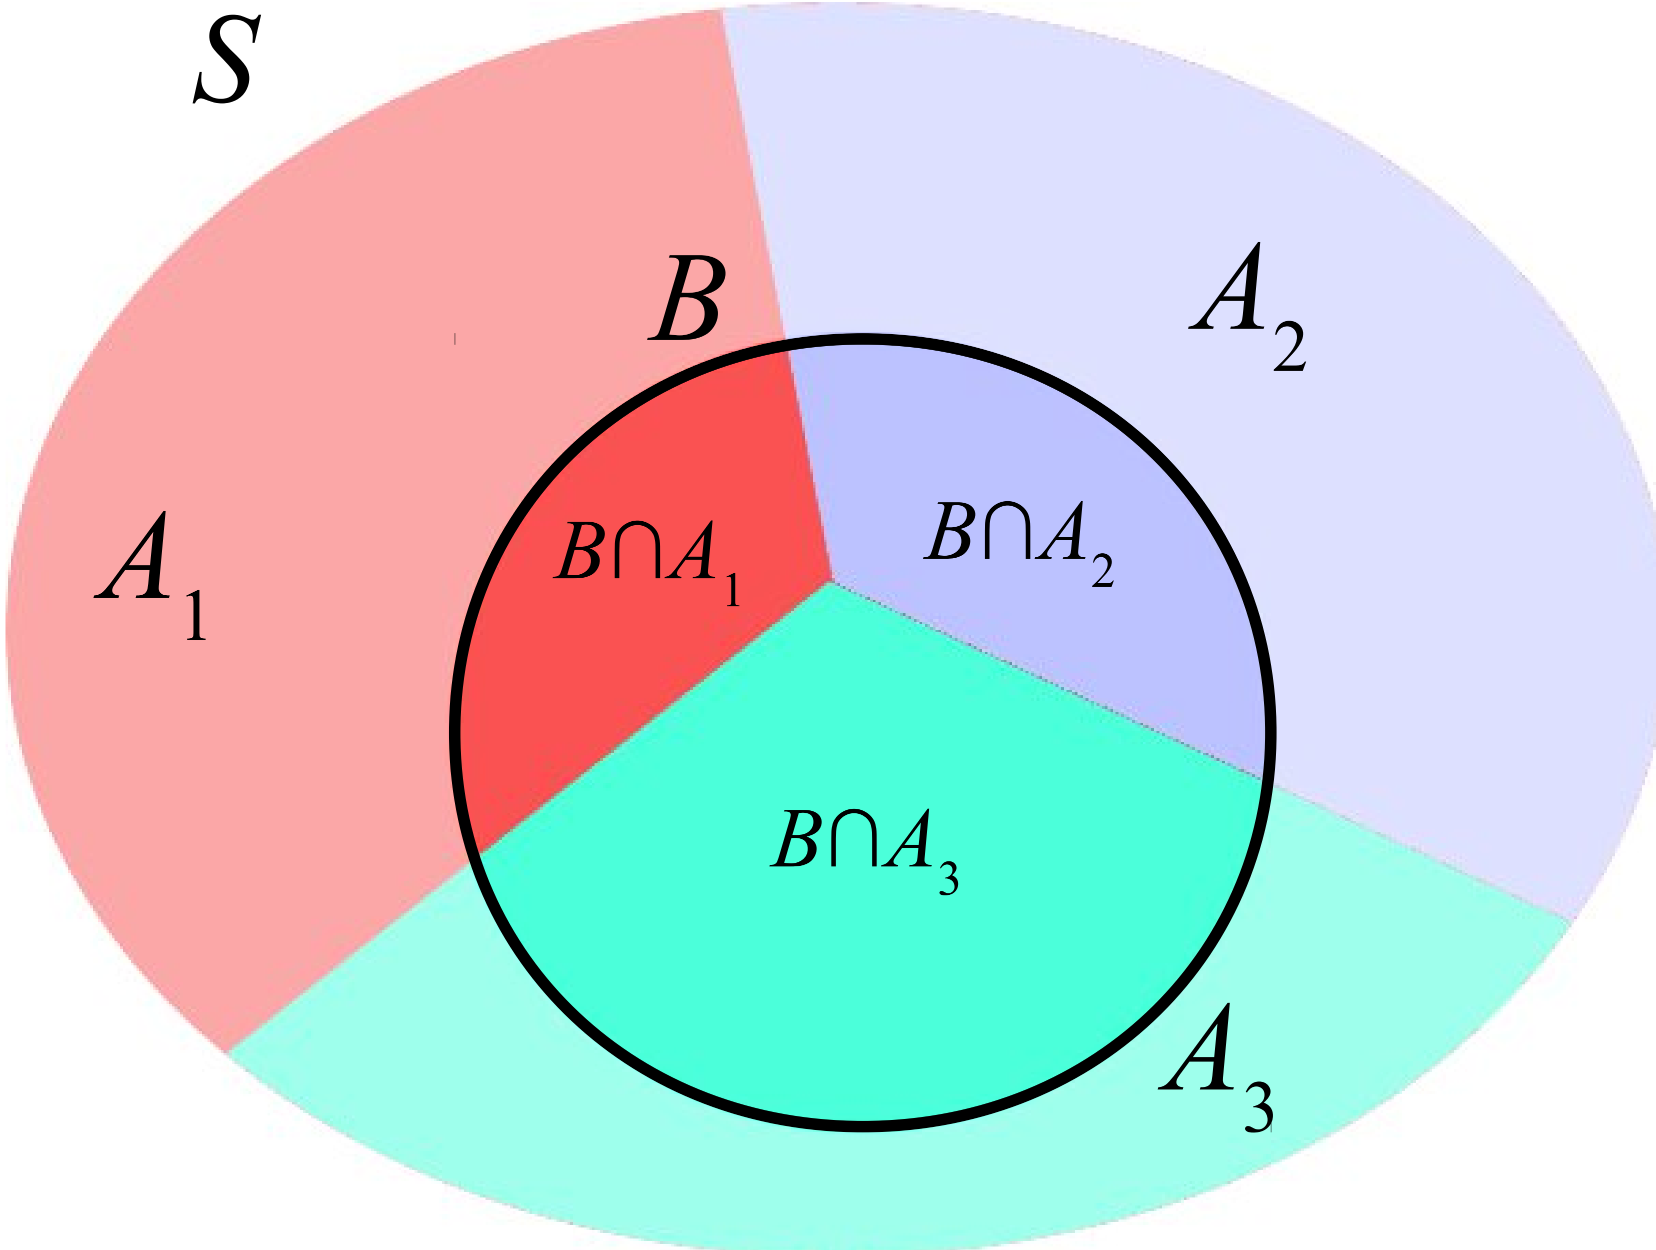
S
B
A_1
A_2
A_3
$B \cap A_1$
$B \cap A_2$
$B \cap A_3$

Recall: Medical Diagnosis

- Suppose:
  - **1 in 1000** people carry a disease, for which there is a pretty reliable test
  - Probability of a false negative (carrier tests negative) is **1%** (so probability of carrier testing positive is **99%**)
  - Probability of a false positive (non-carrier tests positive) is **5%**
- A person just tested positive. What are the chances (s)he is a carrier of the disease?

A counter-intuitive result

- The reliable test gives a positive result, but the chance the patient has the disease is only about **2%**.
- Informally: the rarity of the disease outweighs the small chance of error.

Should we stop worrying?

Scrubs

Should we stop worrying?

- No.

Should we stop worrying?

- No.
- The consequences of being wrong are pretty severe.

Should we stop worrying?

- No.
- The consequences of being wrong are pretty severe.
- How can we better interpret this result?

Possibilities

- You live

Possibilities

- You live

- You die
  (let's assume the disease is always fatal)

Possibilities

- You live
- Improvement to your current state: $+1$

  (if we agree living longer can enrich our life experience)

- You die

  (let's assume the disease is always fatal)

Possibilities

- You live

- Improvement to your current state: +1

  (if we agree living longer can enrich our life experience)

- You die

  (let's assume the disease is always fatal)

- Improvement to your current state: -1000000

Possibilities

- You live
- Improvement to your current state: +1

(if we agree living longer can enrich our life experience)

- You die

(let's assume the disease is always fatal)

- Improvement to your current state: -1000000

(Yes, I'm pulling these numbers out of a (non-probabilistic) posterior, but they're illustrative)

Expectation

$$
\begin{aligned}
&P(\textit{Living} \mid \textit{Positive}) \times \text{Utility of living} \\
&+ P(\textit{Dying} \mid \textit{Positive}) \times \text{Utility of dying}
\end{aligned}
$$

Expectation

$$\begin{aligned} &\mathrm{P}(\textit{Living} \mid \textit{Positive}) \times \text{Utility of living} \\ +\ &\mathrm{P}(\textit{Dying} \mid \textit{Positive}) \times \text{Utility of dying} \end{aligned}$$

$$= 0.986 \times 1 + 0.0194 \times -1000000$$

Expectation

$$\begin{aligned}&P(\textit{Living} \mid \textit{Positive}) \times \text{Utility of living}\\ +\ &P(\textit{Dying} \mid \textit{Positive}) \times \text{Utility of dying}\end{aligned}$$

$$= 0.986 \times 1 + 0.0194 \times -1000000$$

$$\approx -19399$$

Expectation

$$
\begin{aligned}
&\text{P}(\textit{Living} \mid \textit{Positive}) \times \text{Utility of living} \\
&+ \text{P}(\textit{Dying} \mid \textit{Positive}) \times \text{Utility of dying} \\
&= 0.986 \times 1 + 0.0194 \times \text{-}1000000 \\
&\approx \text{-}19399
\end{aligned}
$$

The exact value isn't meaningful in this example. Just note that despite a small probability of a bad outcome, the highly negative weighted average suggests we might have reason to be worried.

Expectation

- The **expectation** (or expected outcome, or expected value…) of an experiment with possible outcomes $\{x_1, x_2, x_3, \ldots\}$ is

$$E(x) = \sum_i P(x_i)\, x_i$$

Expectation

- The **expectation** (or expected outcome, or expected value…) of an experiment with possible outcomes $\{x_1, x_2, x_3, \ldots\}$ is

$$E(x) = \sum_i P(x_i)\, x_i$$

- If the outcomes are not numbers (or more generally, not elements of a *vector space*) which can be meaningfully added/scaled, we can assign them numeric "utility" values $u(x)$, and inspect the "expected utility" $\sum_i P(x_i)\, u(x_i)$

Expected Weather in Ithaca

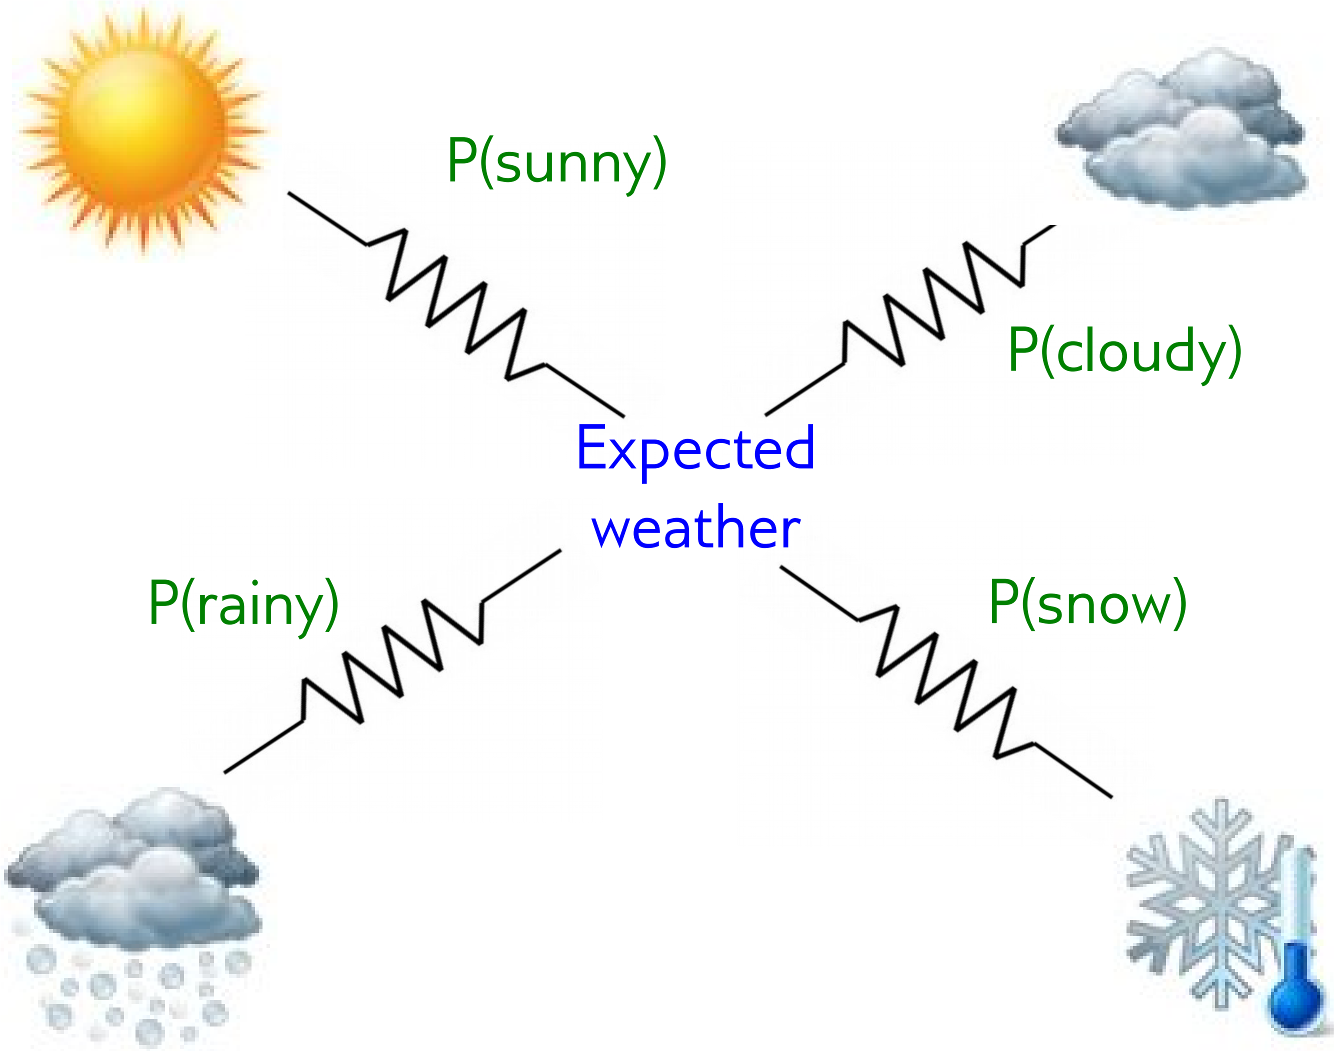

Expected Weather in SoCal

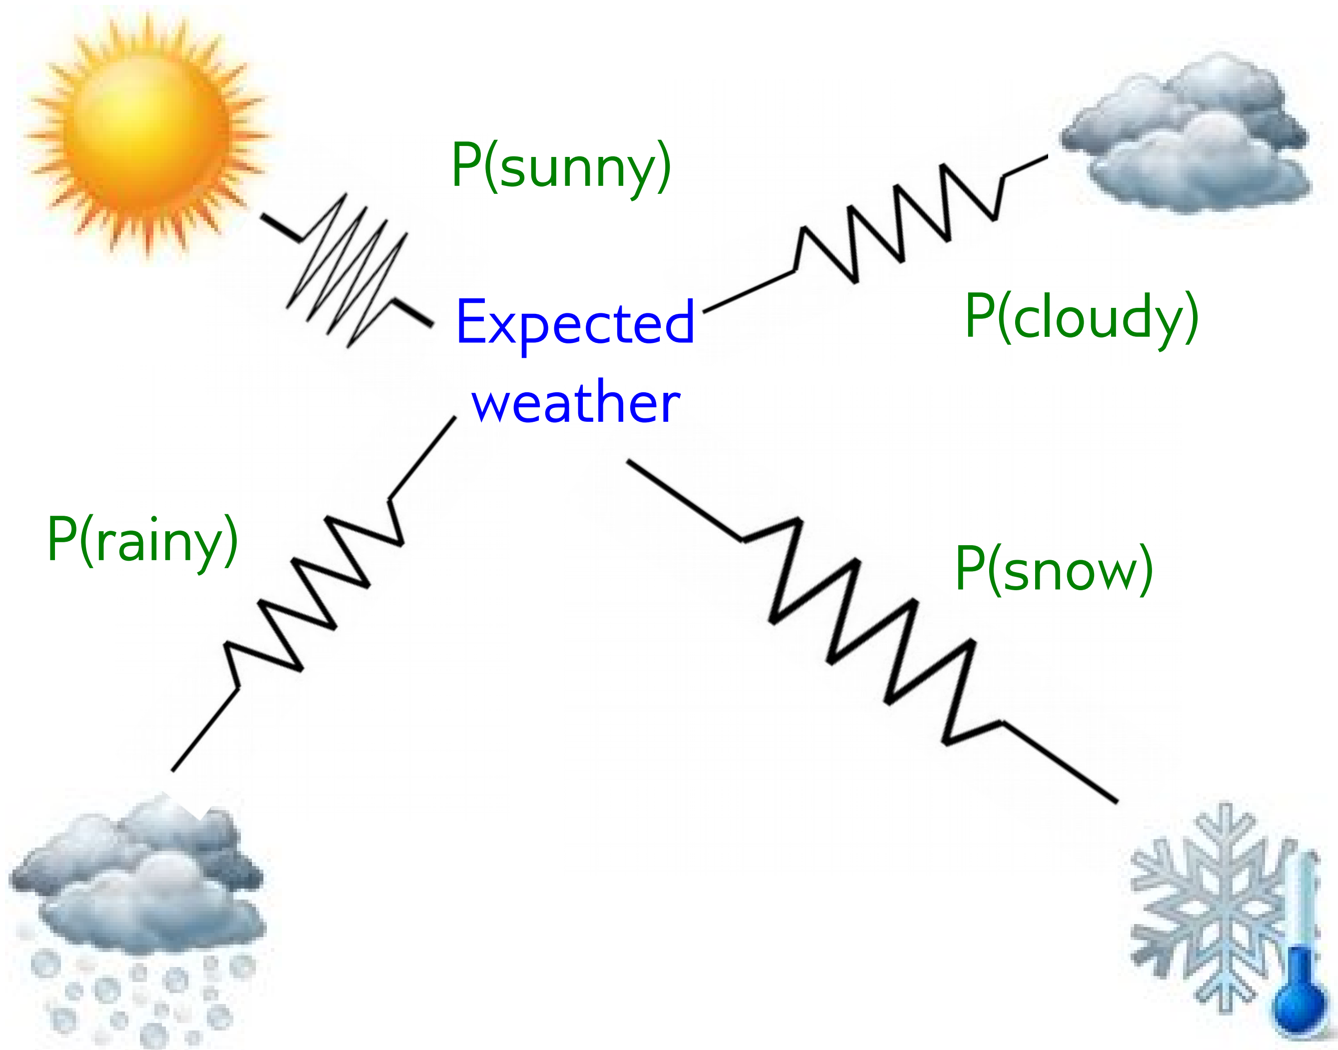

Expectation

- Think of it as a center of mass, with the probabilities being the masses and the outcomes/utilities their (point) positions

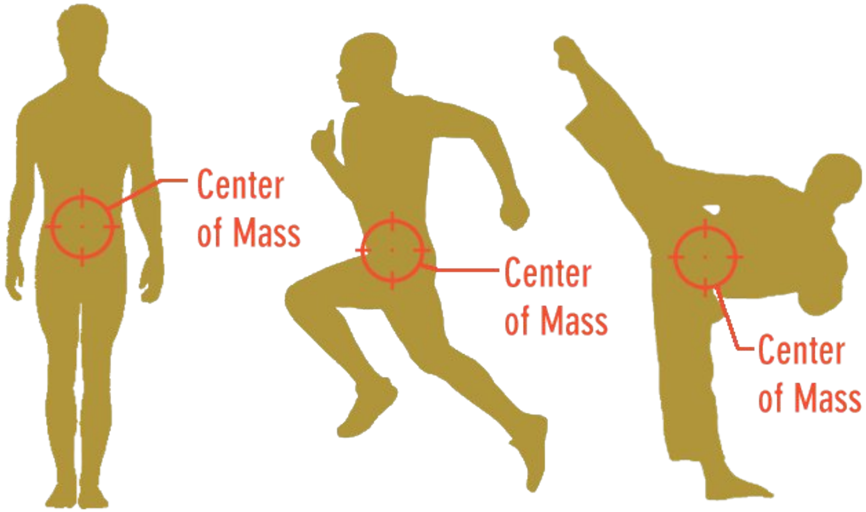

- A useful statistic, but only a guide, not a complete picture of human decision-making

Gold Medal Bodies


MATT GROENING


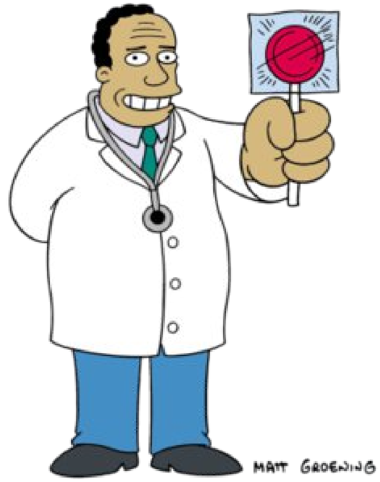
MATT GROENING


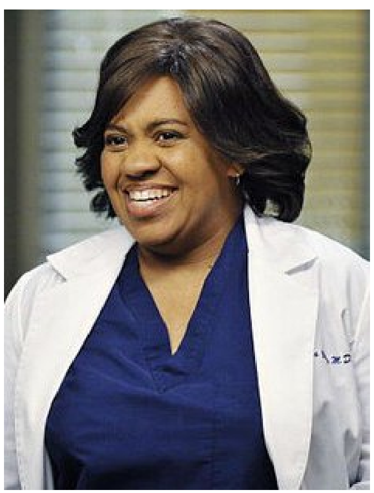

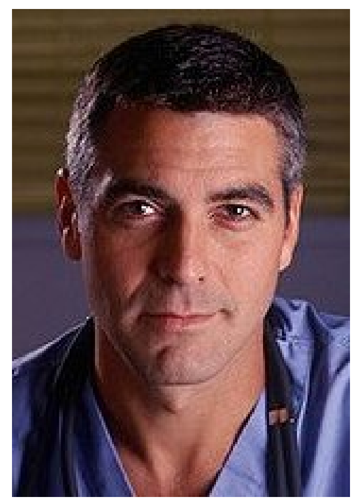


Multiple Tests

- First test: $T_1 = \{ Positive, Negative \}$
- Second test: $T_2 = \{ Positive, Negative \}$

Ignoring patient status for the moment

Multiple Tests

- First test: $T_1 = \{ \mathit{Positive}, \mathit{Negative} \}$
- Second test: $T_2 = \{ \mathit{Positive}, \mathit{Negative} \}$

Ignoring patient status for the moment

- $T = T_1 \times T_2$
  $= \{ (\mathit{Positive}, \mathit{Positive}), (\mathit{Positive}, \mathit{Negative}), (\mathit{Negative}, \mathit{Positive}), (\mathit{Negative}, \mathit{Negative}) \}$

Multiple Tests

- First test: $T_1 = \{ Positive, Negative \}$
- Second test: $T_2 = \{ Positive, Negative \}$

Ignoring patient status for the moment

- $T = T_1 \times T_2$
  $= \{ (Positive, Positive), (Positive, Negative), (Negative, Positive), (Negative, Negative) \}$
- Positive in first test = A_1
  $= ?$

Multiple Tests

- First test: $T_1 = \{ \textit{Positive}, \textit{Negative} \}$
- Second test: $T_2 = \{ \textit{Positive}, \textit{Negative} \}$

Ignoring patient status for the moment

- $T = T_1 \times T_2$
  $= \{ (\textit{Positive}, \textit{Positive}), (\textit{Positive}, \textit{Negative}), (\textit{Negative}, \textit{Positive}), (\textit{Negative}, \textit{Negative}) \}$
- Positive in first test = A_1
  $= \{ (\textit{Positive}, \textit{Positive}), (\textit{Positive}, \textit{Negative}) \}$

Multiple Tests

- First test: $T_1 = \{ Positive, Negative \}$
- Second test: $T_2 = \{ Positive, Negative \}$

Ignoring patient status for the moment

- $T = T_1 \times T_2$

  $= \{ (Positive, Positive), (Positive, Negative), (Negative, Positive), (Negative, Negative) \}$

- Positive in first test = A_1

  $= \{ (Positive, Positive), (Positive, Negative) \}$

- Positive in second test = A_2

  $= ?$

Multiple Tests

- First test: $T_1 = \{ Positive, Negative \}$
- Second test: $T_2 = \{ Positive, Negative \}$

Ignoring patient status for the moment

- $T = T_1 \times T_2$
  $= \{ (Positive, Positive), (Positive, Negative), (Negative, Positive), (Negative, Negative) \}$
- Positive in first test = A_1
  $= \{ (Positive, Positive), (Positive, Negative) \}$
- Positive in second test = A_2
  $= \{ (Positive, Positive), (Negative, Positive) \}$

Multiple Tests

- $D = \{ \textit{Carrier}, \textit{NotCarrier} \}$
- Sample space of first test $S_1 = T_1 \times D$
- Sample space of second test $S_2 = T_2 \times D$
- Full sample space $S = T_1 \times T_2 \times D$
- E.g. event that first test is positive, second is unknown, and person is a carrier:

$$\{ (\boldsymbol{Positive}, Positive, \boldsymbol{Carrier}), (\boldsymbol{Positive}, Negative, \boldsymbol{Carrier}) \}$$

Multiple Tests

- E.g. event that person is a carrier:

{ (*Positive*, *Positive*, ***Carrier***),
(*Positive*, *Negative*, ***Carrier***),
(*Negative*, *Positive*, ***Carrier***),
(*Negative*, *Negative*, ***Carrier***) }

- E.g. event that first test is positive:

{ (***Positive***, *Positive*, *Carrier*),
(***Positive***, *Negative*, *Carrier*),
(***Positive***, *Positive*, *NotCarrier*),
(***Positive***, *Negative*, *NotCarrier*) }

Multiple **Independent** Tests

- We often make the assumption that the test results are independent of each other (but not, of course, of the patient's carrier status). In other words:

$$P(\textit{Test 1 positive} \text{ and } \textit{Test 2 positive} \mid \textit{Carrier})$$
$$= P(\textit{Test 1 positive} \mid \textit{Carrier})$$
$$P(\textit{Test 2 positive} \mid \textit{Carrier})$$

and similarly for all other combinations.

Thought for the Day #1

When might this assumption be unjustified?

Multiple **Independent** Tests

$P\left(C|Y_1 \cap Y_2\right)$

Multiple **Independent** Tests

$$P\left(C \middle| Y_1 \cap Y_2\right)$$

Second test positive
(said "Yes")

Is Carrier

First test positive
(said "Yes")

Multiple **Independent** Tests

$$P\left(C|Y_1 \cap Y_2\right) = \frac{P\left(Y_1 \cap Y_2|C\right) P\left(C\right)}{P\left(Y_1 \cap Y_2\right)}$$

(Bayes' Theorem)

Multiple **Independent** Tests

$$P(C|Y_1 \cap Y_2) = \frac{P(Y_1 \cap Y_2|C)\,P(C)}{P(Y_1 \cap Y_2)}$$

(Bayes' Theorem)

$$= \frac{P(Y_1 \cap Y_2|C)\,P(C)}{P(Y_1 \cap Y_2|C)\,P(C) + P(Y_1 \cap Y_2|C')\,P(C')}$$

(Total Probability Theorem)

Multiple **Independent** Tests

$$P\left(C|Y_1 \cap Y_2\right) = \frac{P\left(Y_1 \cap Y_2|C\right)P\left(C\right)}{P\left(Y_1 \cap Y_2\right)}$$

(Bayes' Theorem)

$$= \frac{P\left(Y_1 \cap Y_2|C\right)P\left(C\right)}{P\left(Y_1 \cap Y_2|C\right)P\left(C\right) + P\left(Y_1 \cap Y_2|C'\right)P\left(C'\right)}$$

(Total Probability Theorem)

$$= \frac{P\left(Y_1|C\right)P\left(Y_2|C\right)P\left(C\right)}{P\left(Y_1|C\right)P\left(Y_2|C\right)P\left(C\right) + P\left(Y_1|C'\right)P\left(Y_2|C'\right)P\left(C'\right)}$$

(Independence assumption)

Multiple **Independent** Tests

$$P(C|Y_1 \cap Y_2) = \frac{P(Y_1 \cap Y_2|C)\,P(C)}{P(Y_1 \cap Y_2)}$$

(Bayes' Theorem)

$$= \frac{P(Y_1 \cap Y_2|C)\,P(C)}{P(Y_1 \cap Y_2|C)\,P(C) + P(Y_1 \cap Y_2|C')\,P(C')}$$

(Total Probability Theorem)

$$= \frac{P(Y_1|C)\,P(Y_2|C)\,P(C)}{P(Y_1|C)\,P(Y_2|C)\,P(C) + P(Y_1|C')\,P(Y_2|C')\,P(C')}$$

(Independence assumption)

$$= \frac{0.99 \times 0.99 \times 0.001}{0.99 \times 0.99 \times 0.001 + 0.05 \times 0.05 \times 0.999}$$

(From problem statement. The modeling assumption is that these values also apply to the product space)

Multiple **Independent** Tests

$$P(C|Y_1 \cap Y_2) = \frac{P(Y_1 \cap Y_2|C)\,P(C)}{P(Y_1 \cap Y_2)}$$

(Bayes' Theorem)

$$= \frac{P(Y_1 \cap Y_2|C)\,P(C)}{P(Y_1 \cap Y_2|C)\,P(C) + P(Y_1 \cap Y_2|C')\,P(C')}$$

(Total Probability Theorem)

$$= \frac{P(Y_1|C)\,P(Y_2|C)\,P(C)}{P(Y_1|C)\,P(Y_2|C)\,P(C) + P(Y_1|C')\,P(Y_2|C')\,P(C')}$$

(Independence assumption)

$$= \frac{0.99 \times 0.99 \times 0.001}{0.99 \times 0.99 \times 0.001 + 0.05 \times 0.05 \times 0.999}$$

(From problem statement. The modeling assumption is that these values also apply to the product space)

$$= \mathbf{0.2818}$$

Thought for the Day #2

Can you write down, in full, all the outcomes in each event involved in the previous slide? Note that these events are all subsets of $S = T_1 \times T_2 \times D$

Thought for the Day #3

Can you work out the probability of the patient being a carrier, for every other combination of test results?