Preference Modeling

Lirong Xia

Rensselaer

Today's Schedule

- Modeling random preferences
  - Random Utility Model
- Modeling preferences over lotteries
  - Prospect theory

Parametric Preference Learning

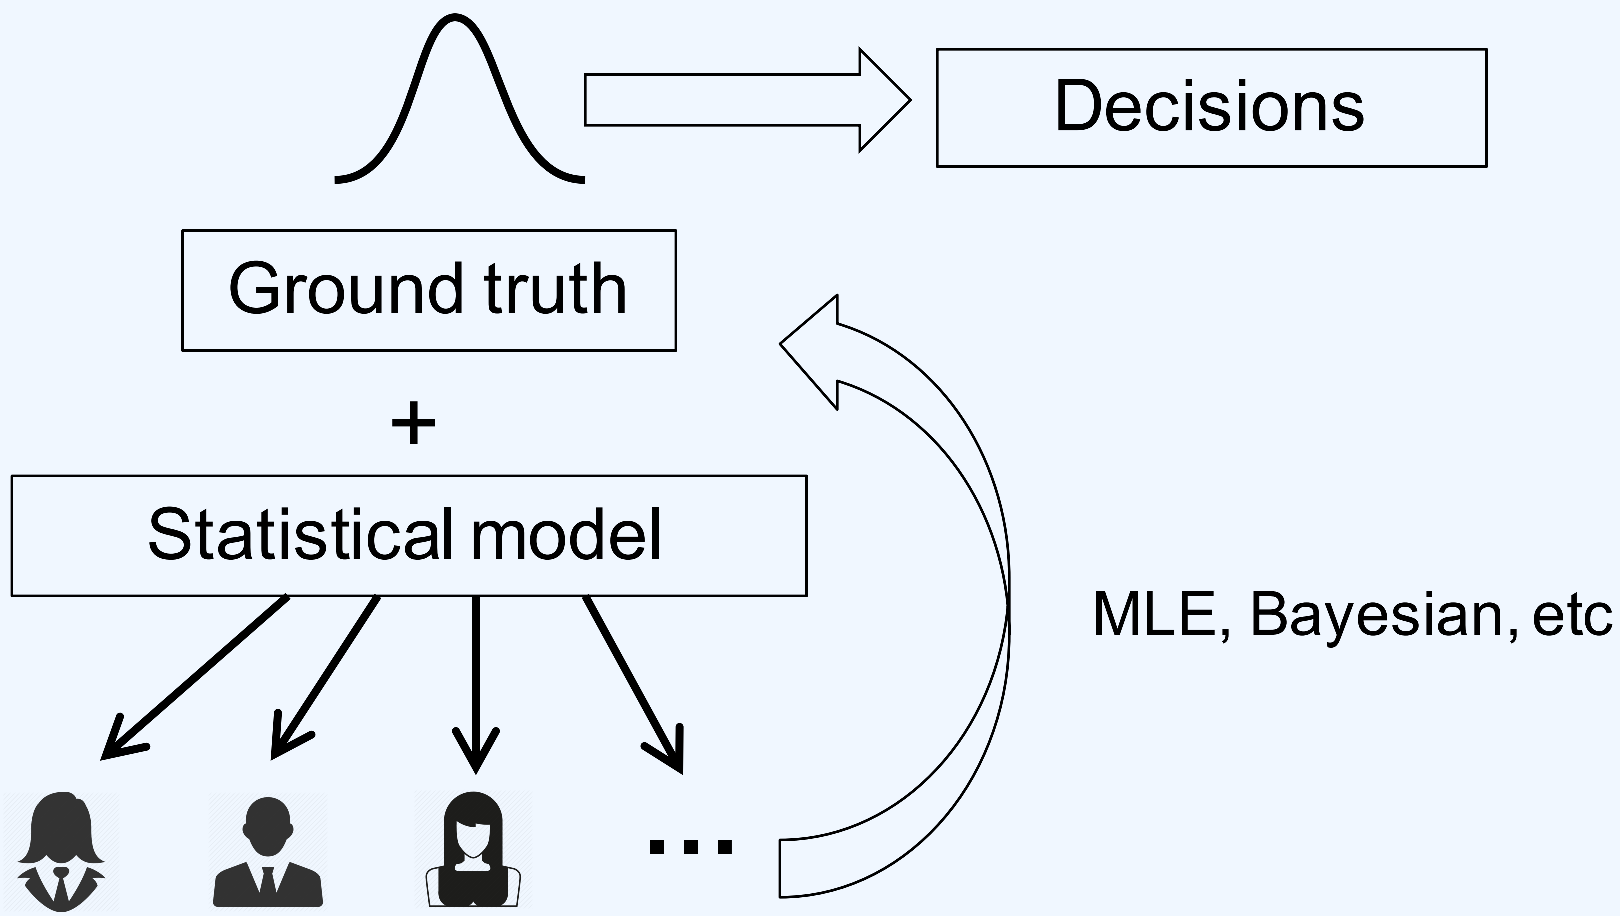

Parametric ranking models

- A statistical model has three parts
  - A parameter space: Θ
  - A sample space: S = Rankings(A)^n
    - A = the set of alternatives, n=#voters
    - assuming votes are i.i.d.
  - A set of probability distributions over S:

    $\{Pr_\theta(s)$ for each $s \in$ Rankings(A) and $\theta \in \Theta\}$

Example

- Condorcet's model for two alternatives
- Parameter space Θ={  , 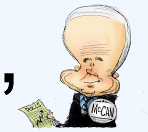 }

- Sample space S = {  , 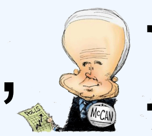 }^{n}

- Probability distributions, i.i.d.

Pr( | )

= Pr(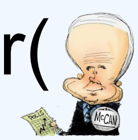 | )

= $p > 0.5$

Mallows' model [Mallows-1957]

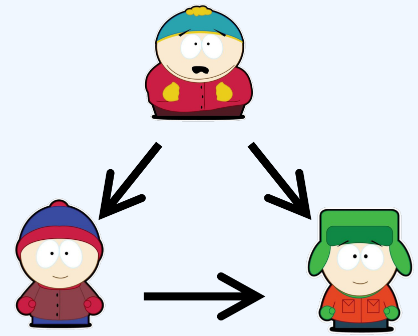

- Fixed dispersion φ <1
- Parameter space
  - all full rankings over candidates
- Sample space
  - i.i.d. generated full rankings
- Probabilities:

$$\Pr_W(V) \propto \varphi^{\text{Kendall}(V,W)}$$

Example: Mallows for

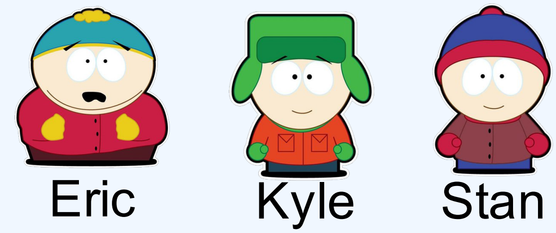

- Probabilities: $Z = 1 + 2\varphi + 2\varphi^2 + \varphi^3$

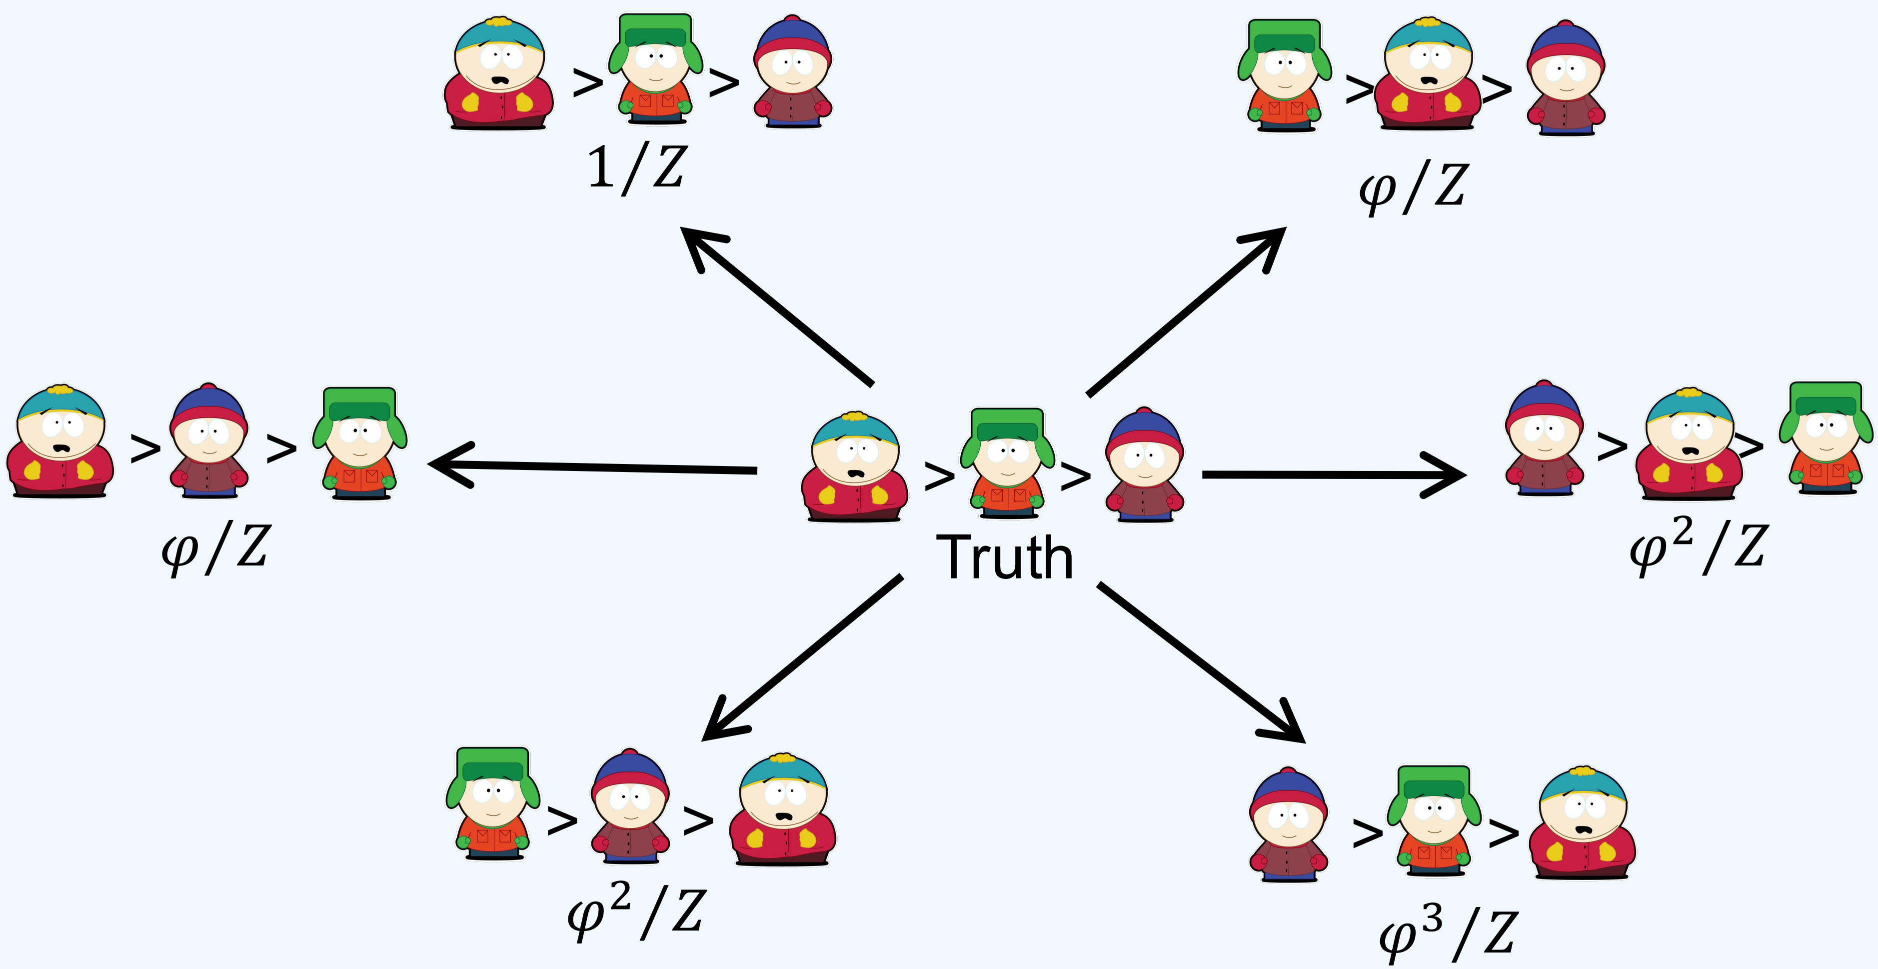

Random utility model (RUM)
[Thurstone 27]

- Continuous parameters: $\Theta=(\theta_1,\ldots,\theta_m)$
  - m: number of alternatives
  - Each alternative is modeled by a utility distribution μ_i
  - θ_i: a vector that parameterizes μ_i
- An agent's latent utility U_i for alternative c_i is generated independently according to $\mu_i(U_i)$
- Agents rank alternatives according to their perceived utilities
  - $\Pr(c_2 \succ c_1 \succ c_3 | \theta_1, \theta_2, \theta_3) = \Pr_{U_i \sim \mu_i}(U_2 > U_1 > U_3)$

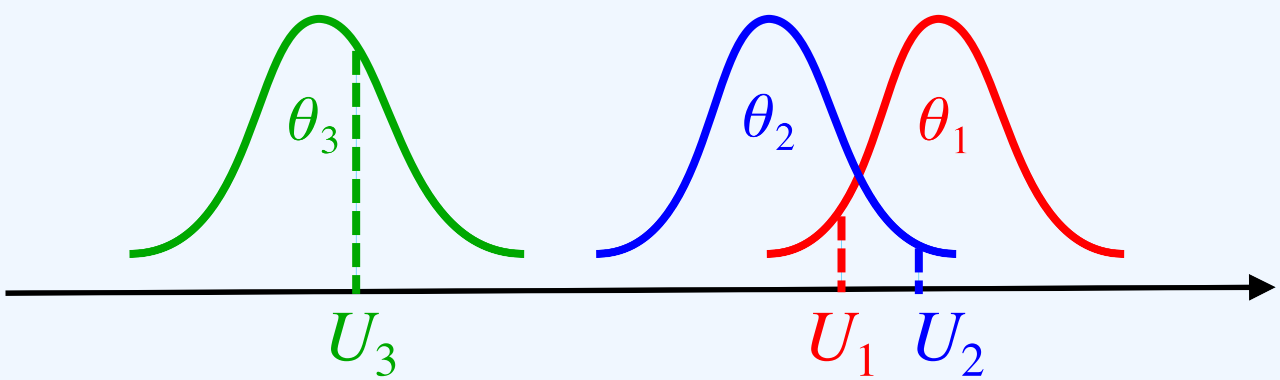

Generating a preference-profile

- $\Pr(\text{Data} \mid \theta_1, \theta_2, \theta_3) = \prod_{V \in \text{Data}} \Pr(V \mid \theta_1, \theta_2, \theta_3)$

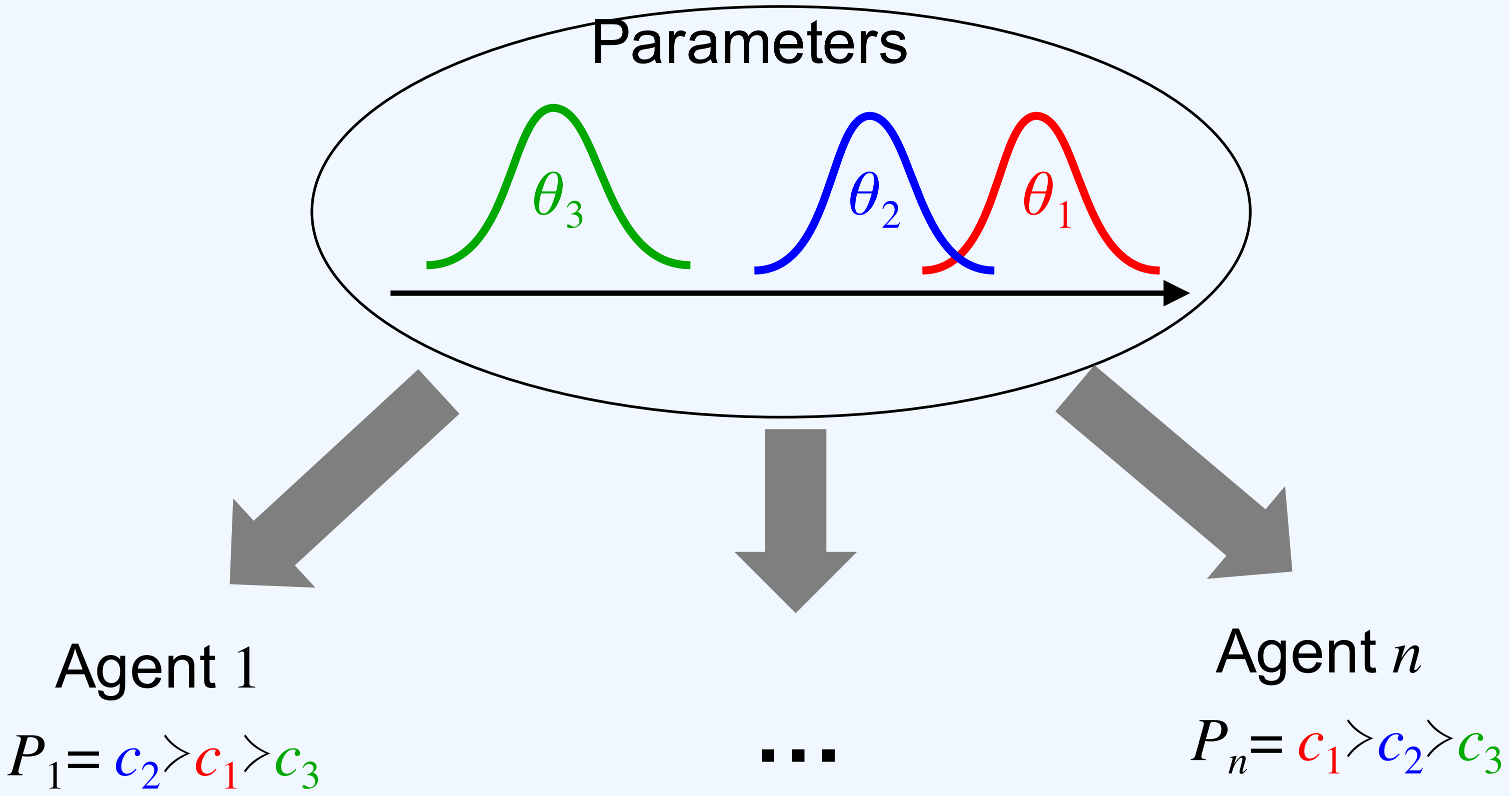

Plackett-Luce model

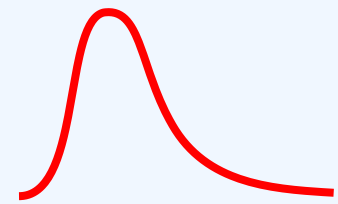

- μ_i's are Gumbel distributions
  - A.k.a. the Plackett-Luce (P-L) model [BM 60, Yellott 77]
- Alternative parameterization $\lambda_1,\ldots,\lambda_m$

$$\Pr(c_1 \succ c_2 \succ \cdots \succ c_m \mid \lambda_1 \cdots \lambda_m) = \frac{\lambda_1}{\lambda_1 + \cdots + \lambda_m} \times \frac{\lambda_2}{\lambda_2 + \cdots + \lambda_m} \times \cdots \times \frac{\lambda_{m-1}}{\lambda_{m-1} + \lambda_m}$$

c_1 is the top choice in $\{c_1,\ldots,c_m\}$
c_2 is the top choice in $\{c_2,\ldots,c_m\}$
c_{m-1} is preferred to c_m

- Pros:
  - Computationally tractable
    - Analytical solution to the likelihood function
      - The only RUM that was known to be tractable
    - Widely applied in Economics [McFadden 74], learning to rank [Liu 11], and analyzing elections [GM 06,07,08,09]

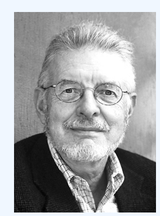

McFadden

- Cons: may not be the best model

Example

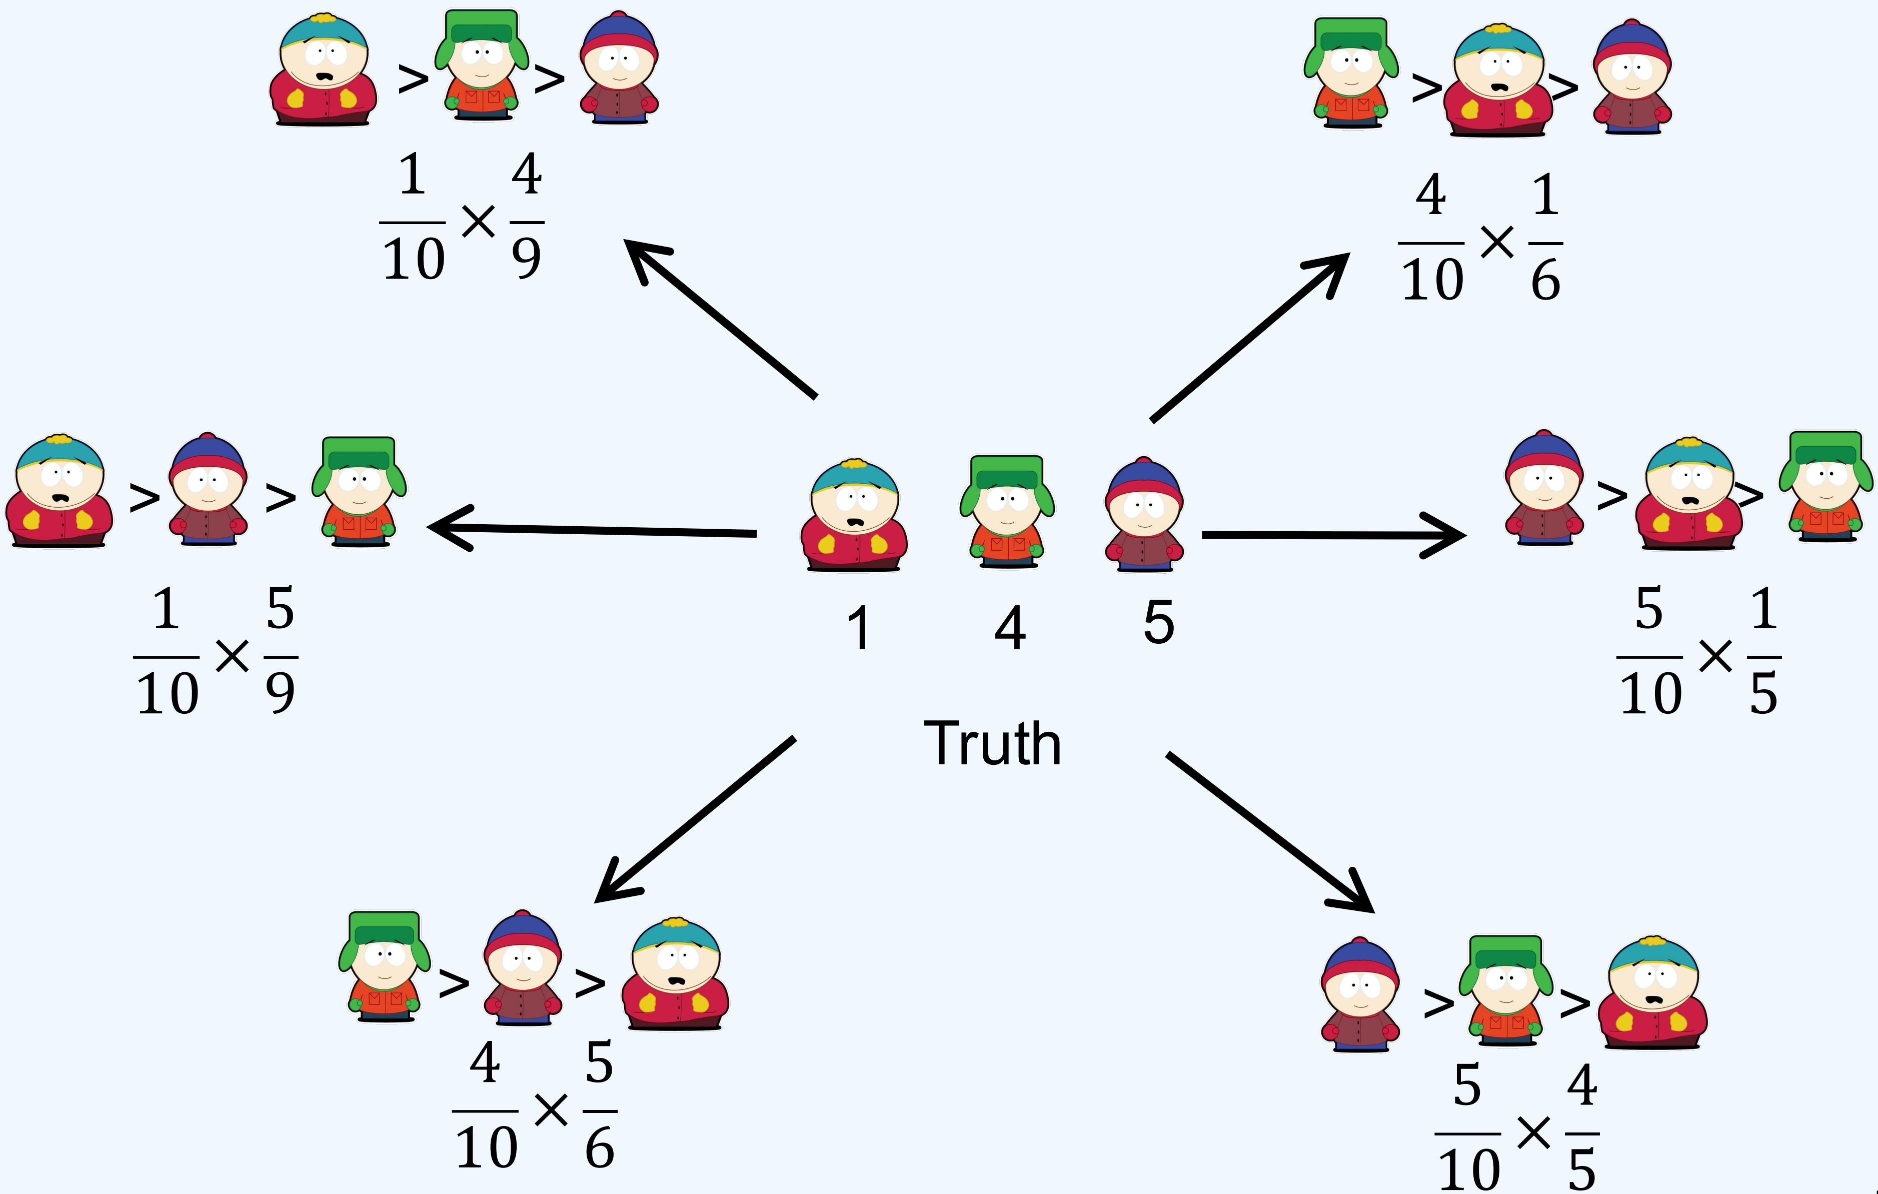
$\frac{1}{10} \times \frac{4}{9}$
$\frac{4}{10} \times \frac{1}{6}$
$\frac{1}{10} \times \frac{5}{9}$
1 4 5
Truth
$\frac{5}{10} \times \frac{1}{5}$
$\frac{4}{10} \times \frac{5}{6}$
$\frac{5}{10} \times \frac{4}{5}$

RUM with normal distributions

- μ_i's are normal distributions
  - Thurstone's Case V [Thurstone 27]

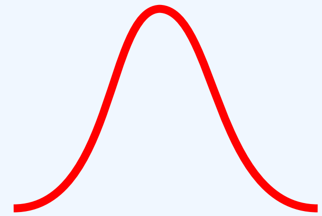

- Pros:
  - Intuitive
  - Flexible
- Cons: believed to be computationally intractable
  - No analytical solution for the likelihood function $\Pr(P \mid \Theta)$ is known

$$\Pr(c_1 \succ \cdots \succ c_m \mid \Theta) = \int_{-\infty}^{\infty} \int_{U_m}^{\infty} \cdots \int_{U_2}^{\infty} \mu_m(U_m)\, \mu_{m-1}(U_{m-1}) \cdots \mu_1(U_1)\, dU_1 \cdots dU_{m-1}\, dU_m$$

U_m: from -∞ to ∞ | U_{m-1}: from U_m to ∞ | ⋯ | U_1: from U_2 to ∞

Decision making

Maximum likelihood estimators (MLE)

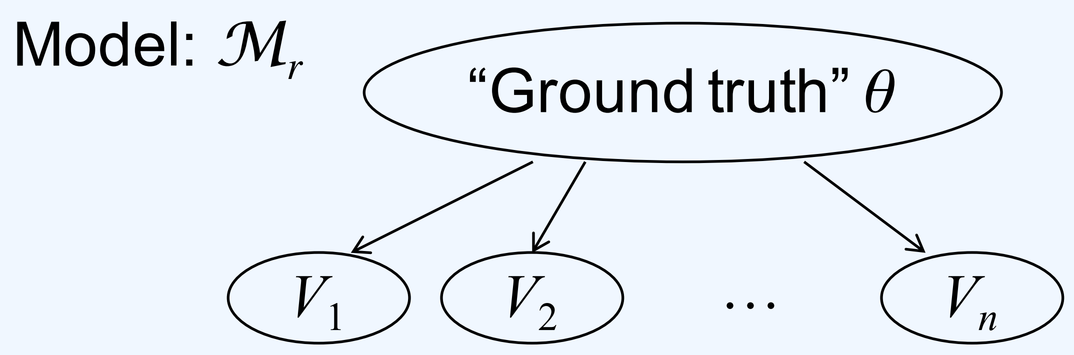

➢ For any profile $P=(V_1,\ldots,V_n)$,

- The likelihood of θ is $L(\theta,P)=\Pr_\theta(P)=\prod_{V\in P}\Pr_\theta(V)$
- The MLE mechanism

  MLE$(P)=$argmax$_\theta\, L(\theta,P)$
- Decision space = Parameter space

Bayesian approach

- Given a profile $P=(V_1,\ldots,V_n)$, and a prior distribution π over Θ
- Step 1: calculate the posterior probability over Θ using Bayes' rule
  - $\Pr(\theta|P) \propto \pi(\theta) \Pr_\theta(P)$
- Step 2: make a decision based on the posterior distribution
  - Maximum a posteriori (MAP) estimation
  - $\text{MAP}(P)=\text{argmax}_\theta \Pr(\theta|P)$
  - Technically equivalent to MLE when π is uniform

Example

- Θ={ , 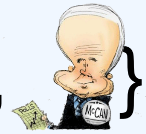}
- S = { , 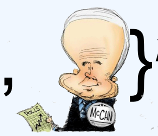}^{n}
- Probability distributions:
- Data P = {10@  + 8@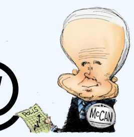 }
- MLE
  - $L(O)=Pr_O(O)^6\ Pr_O(M)^4 = 0.6^{10}\ 0.4^8$
  - $L(M)=Pr_M(O)^6\ Pr_M(M)^4 = 0.4^{10}\ 0.6^8$
  - L(O)>L(M), O wins
- MAP: prior O:0.2, M:0.8
  - $Pr(O|P) \propto 0.2\ L(O) = 0.2 \times 0.6^{10}\ 0.4^8$
  - $Pr(M|P) \propto 0.8\ L(M) = 0.8 \times 0.4^{10}\ 0.6^8$
  - Pr(M|P)> Pr(O|P), M wins

Pr( | )

= Pr(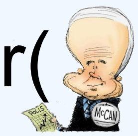 | 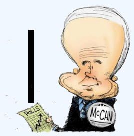)

$= 0.6$

Decision making under uncertainty

- You have a biased coin: head w/p p
  - You observe 10 heads, 4 tails
  - Do you think the next two tosses will be two heads in a row?

Credit: Panos Ipeirotis & Roy Radner

- MLE-based approach
  - there is an unknown but fixed ground truth
  - p = 10/14=0.714
  - Pr(2heads|p=0.714) =$(0.714)^2$=0.51>0.5
  - Yes!

- Bayesian
  - the ground truth is captured by a belief distribution
  - Compute Pr(p|Data) assuming uniform prior
  - Compute Pr(2heads|Data)=0.485<0.5
  - No!

Prospect Theory: Motivating Example

- Treat lung cancer with Radiation or Surgery

	Radiation	Surgery
Q1	100% immediately survive 22% 5-year survive	90% immediately survive 34% 5-year survive

- Q1: 18% choose Radiation
- Q2: 49% choose Radiation

More Thoughts

- Framing Effect
  - The baseline/starting point matters
    - Q1: starting at “the patient dies”
    - Q2: starting at “the patient survives”
- Evaluation
  - subjective value (utility)
  - perceptual likelihood (e.g. people tend to overweight low probability events)

Kahneman

Prospect Theory

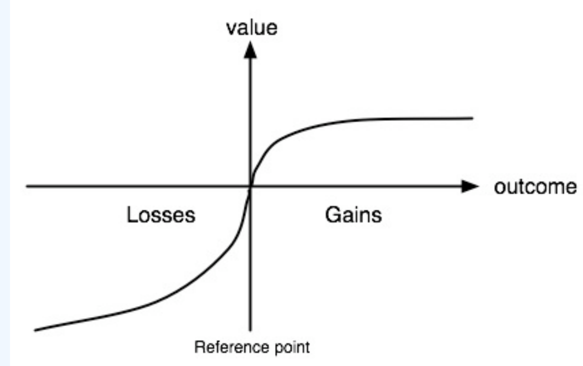

- Framing Phase (modeling the options)
  - Choose a reference point to model the options
  - $O = \{o_1,\ldots,o_k\}$
- Evaluation Phase (modeling the preferences)
  - a value function $v: O \rightarrow R$
  - a probability weighting function $\pi: [0,1] \rightarrow R$
- For any lottery $L= (p_1,\ldots,p_k) \in Lot(O)$

$$V(L) = \sum \pi(pi)v(o_i)$$

Example: Insurance

- potential loss of \$1000 @1%
- Insurance fee \$15
- Q1 (reference point: current wealth)
  - Buy: Pay \$15 for sure. V = v(-15)
  - No: \$0@99% + \$-1000@1%. V = $\pi(.99)v(0) + \pi(.01)v(-1000) = \pi(.01)v(-1000)$
- Q2 (reference point: current wealth-1000)
  - Buy: \$985 for sure. V = v(985)
  - No: \$1000@99% + \$0@1%. V = $\pi(.99)v(1000) + \pi(.01)v(0) = \pi(.99)v(1000)$