CPS 196.2
Proper scoring rules

Vincent Conitzer
conitzer@cs.duke.edu

How to incentivize a weather forecaster

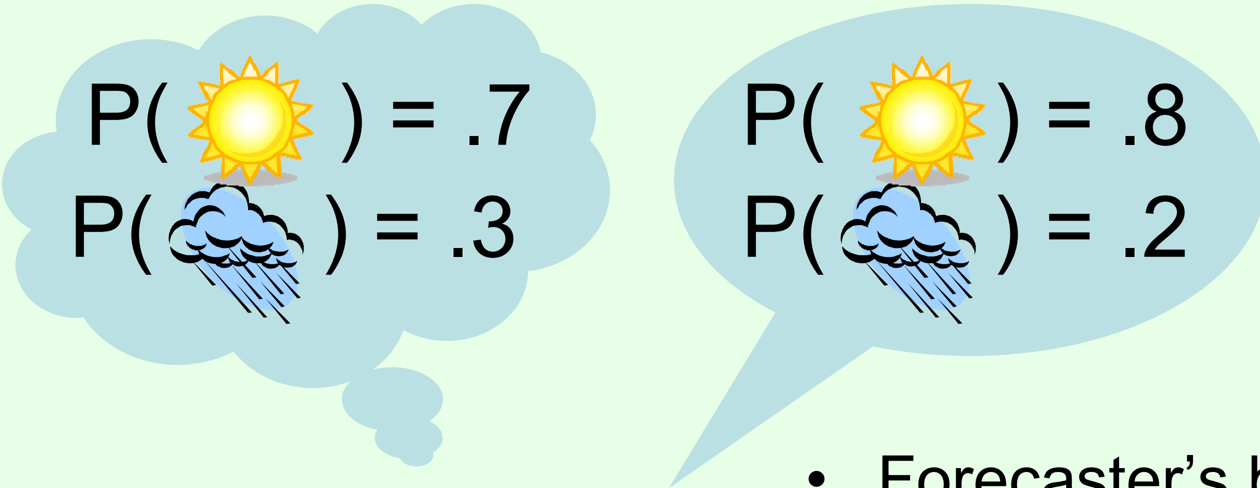

- Forecaster’s bonus can depend on
  - Prediction
  - Actual weather on predicted day
- Reporting true beliefs should maximize expected bonus

A first attempt

- Suppose we reward the forecaster as follows:
  - Suppose outcome i happened
  - Pay the forecaster the probability he assigned to i
- Notation (in binary setting):
  - x = whether the event (say, rain) happens
    - 1 if rain, 0 if sun
  - q = probability (of rain) reported by forecaster
  - p = forecaster's private (true) probability (of rain)
- So, forecaster receives

  $xq + (1-x)(1-q)$
- Forecaster's expected payoff for reporting q:

  $pq + (1-p)(1-q)$
- How does forecaster choose q to maximize expected payoff?

A different reward function

- How about: $xq - q^2/2$
- Forecaster's expected payoff: $pq - q^2/2$
- Derivative w.r.t. q: $p - q$
  – Setting to 0 gives $q=p$
  – Note second derivative is negative
- Forecaster is (strictly) incentivized to tell the truth!
- We say that $xq - q^2/2$ is a proper scoring rule
- Little funny: asymmetric between $x=0$ and $x=1$

Brier (aka. quadratic) scoring rule

- How about: $1 - (x - q)^2$
- Forecaster's expected payoff:
- $1 - p(1 - q)^2 - (1 - p)q^2$
- Derivative w.r.t. q: $-2pq + 2p - 2(1-p)q = 2p-2q$
  - Setting to 0 gives $q=p$
  - Note second derivative is negative
- So Brier is also proper

Logarithmic scoring rule

- How about x log q + (1-x) log (1-q)
- Forecaster's expected payoff:
- p log q + (1-p) log (1-q)
- Derivative w.r.t. q: p/q - (1-p)/(1-q)
  - Setting to 0 gives q=p
  - Note second derivative is negative
- So logarithmic is also proper

A securities-based interpretation

- Let's say we allow the forecaster to buy "rain securities"
- Each rain security pays off 1 if it rains, 0 otherwise
- Price starts at 0;
- After the forecaster buys s(0) securities, we increase the price to .1;
- After the forecaster buys s(.1) securities, we increase the price to .2;
- Etc.
- Forecaster should keep buying until price = p
  - (regardless of function s)

A securities-based interpretation

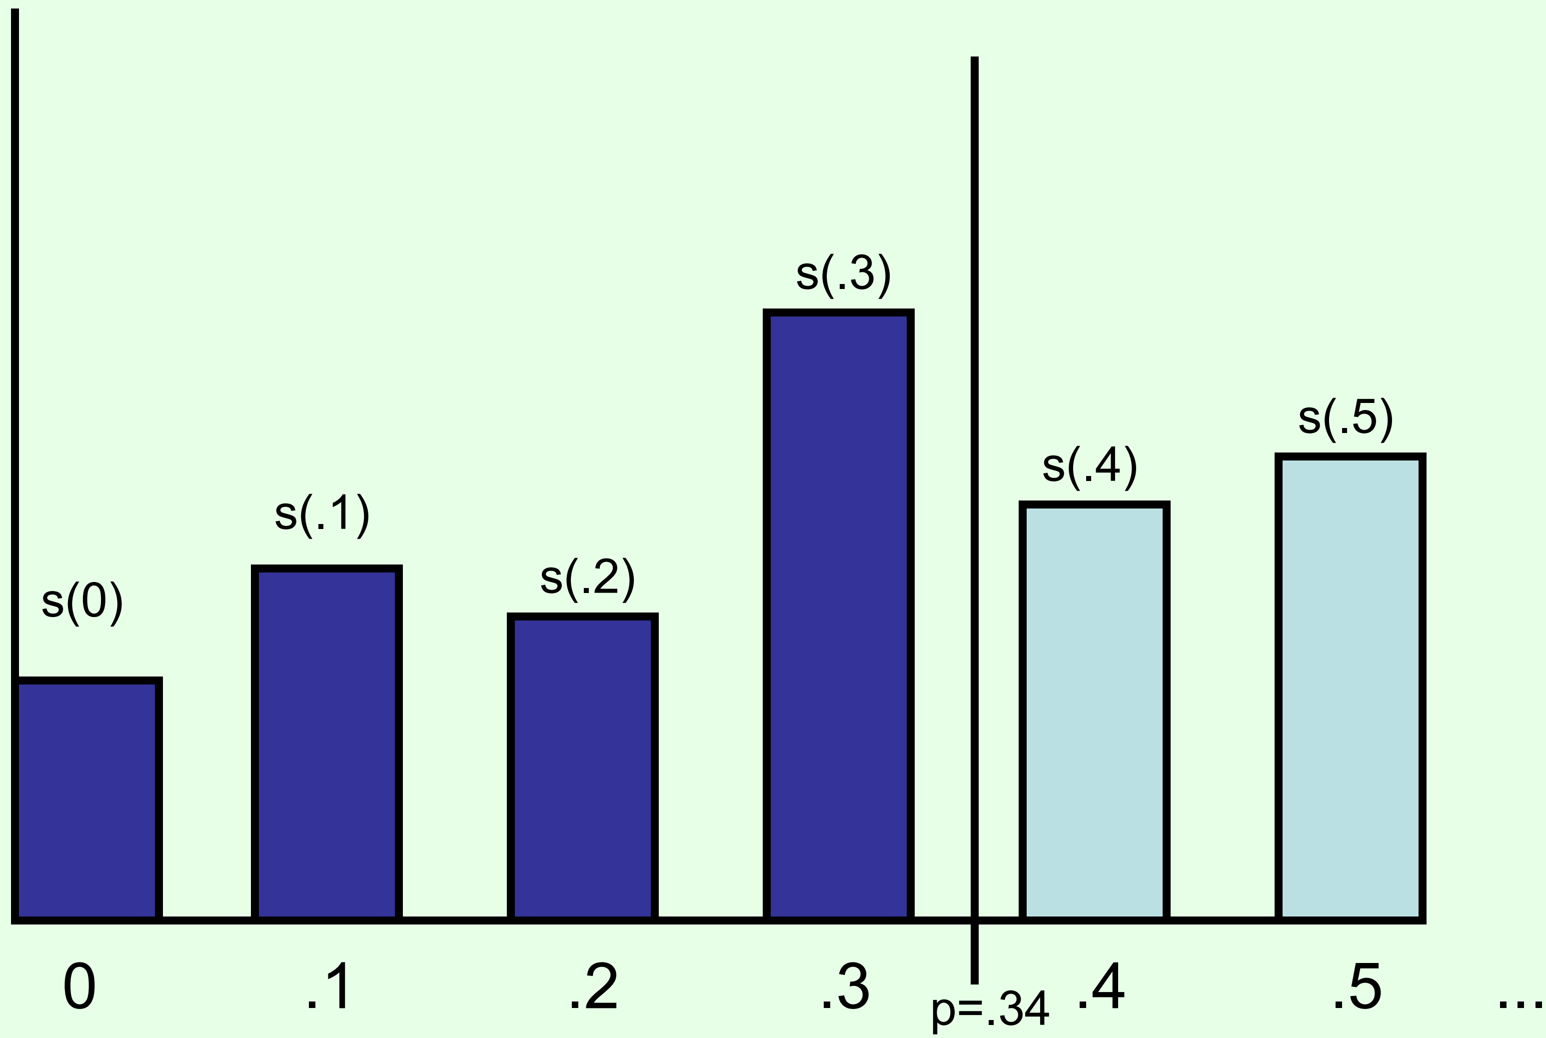

Forecaster pays s(0)*0 + s(.1)*.1 + s(.2)*.2 + s(.3)*.3,
expects .34*(s(0)+s(.1)+s(.2)+s(.3))

Formulas

- If the forecaster buys up to price q, he will end up buying

  $\sum_{y=0}^{q} s(y)$

  securities (y is the price)
- He will pay

  $\sum_{y=0}^{q} s(y)y$
- He expects a payoff of

  $p^* \sum_{y=0}^{q} s(y)$

Making things continuous

- If the forecaster buys up to price q, he will end up buying

  $\int_{y=0}^{q} s(y)dy$

  securities
- He will pay

  $\int_{y=0}^{q} s(y)y\,dy$
- He expects a payoff of

  $p^* \int_{y=0}^{q} s(y)dy$

Example

- $s(y) = 1$
- If the forecaster buys up to price q, he will pay $_{y=0}\int^{q} s(y)y\,dy = {}_{y=0}\int^{q} y\,dy = q^2/2$
- He will receive $x * {}_{y=0}\int^{q} s(y)\,dy = x * {}_{y=0}\int^{q} dy = xq$
- Total payoff = $xq - q^2/2$
- This was the proper scoring rule we started with!

Another example

- $s(y) = y$
- If the forecaster buys up to price q, he will pay $\int_{y=0}^{q} s(y)y\,dy = \int_{y=0}^{q} y^2 dy = q^3/3$
- He will receive $x^* \int_{y=0}^{q} s(y)\,dy = x^* \int_{y=0}^{q} y\,dy = xq^2/2$
- Total payoff $= xq^2/2 - q^3/3$
- Expected payoff $= pq^2/2 - q^3/3$
- Derivative w.r.t. $q = pq - q^2$
  - Expected payoff maximized at $q=p$
- I.e. $xq^2/2 - q^3/3$ is also a proper scoring rule
  - Not surprising, since of course you want to keep buying up to where the price hits p

Yet another example

- $s(y) = e^y$
- If the forecaster buys up to price q, he will pay $\int_{y=0}^{q} s(y)y\,dy = \int_{y=0}^{q} e^y y\,dy = (q-1)e^q$
- He will receive $x^* \int_{y=0}^{q} s(y)dy = x^* \int_{y=0}^{q} e^y dy = xe^q$
- Total payoff = $xe^q - (q-1)e^q$
- Expected payoff = $pe^q - (q-1)e^q$
- Derivative w.r.t. q = $pe^q - (q-1)e^q - e^q$
  - Expected payoff maximized at q=p
- I.e. $xe^q - (q-1)e^q$ is also a proper scoring rule
  - Not surprising, since of course you want to keep buying up to where the price hits p

More than two outcomes

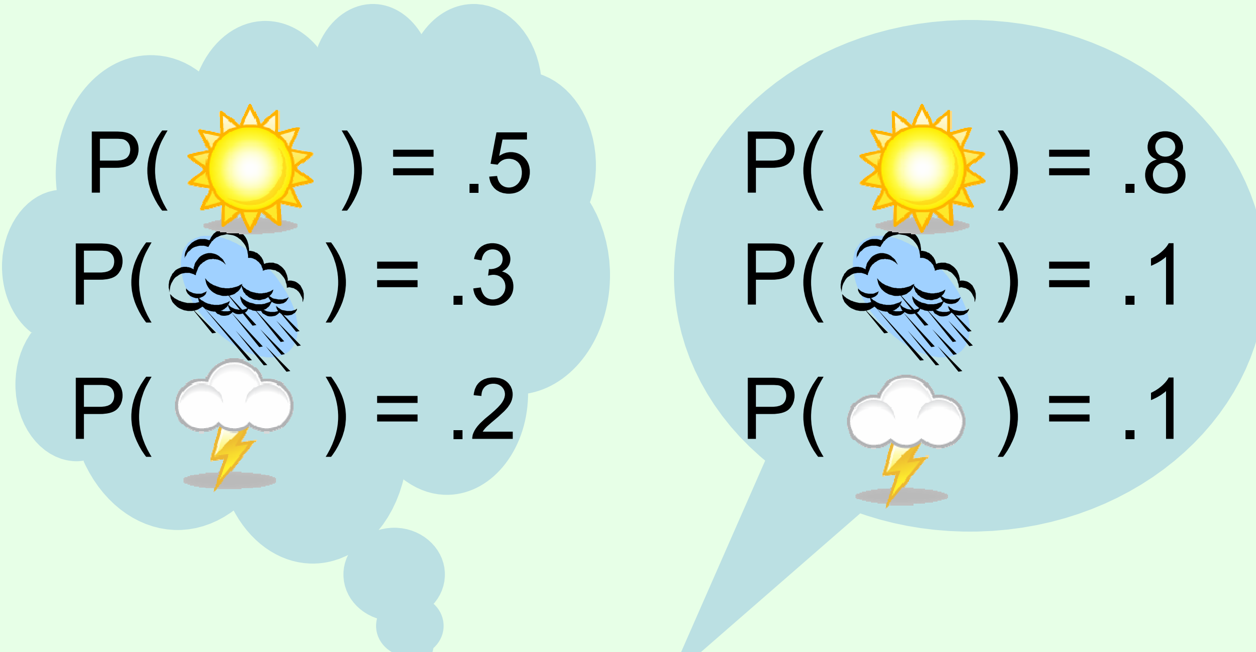

- Can extend all the above by simply asking for separate predictions for whether there will be sun, whether there will be rain, whether there will be lightning

Market scoring rules

- Extend proper scoring rules to multiple forecasters
- Each forecaster is allowed to move current distribution
- Gets rewarded according to new distribution, *minus* reward for previous distribution
- Say we use proper scoring rule $xq - q^2/2$

time

P() = .3

P() = .5

P() = .4

$x(.3) - (.3)^2/2$

$x(.5) - (.5)^2/2 - x(.3) + (.3)^2/2$

$x(.4) - (.4)^2/2 - x(.5) + (.5)^2/2$