Robust probability updating

Thijs van Ommen

UvA

June 7, 2016

Monty Hall's game show

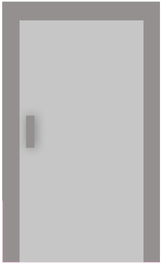 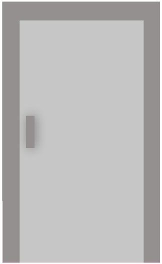 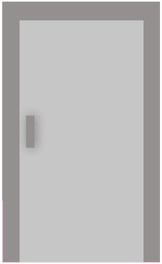

Initial probability:

$1/3$ $1/3$ $1/3$

Monty Hall's game show

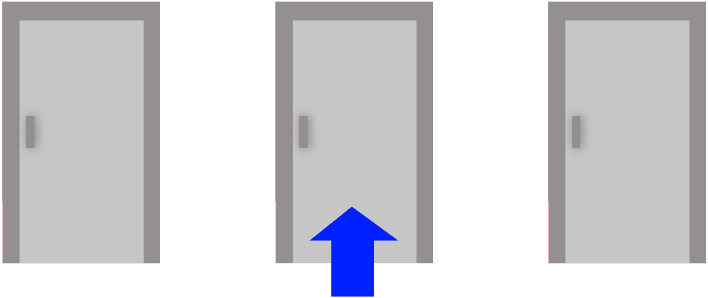

Initial probability:

1/3 1/3 1/3

Monty Hall's game show

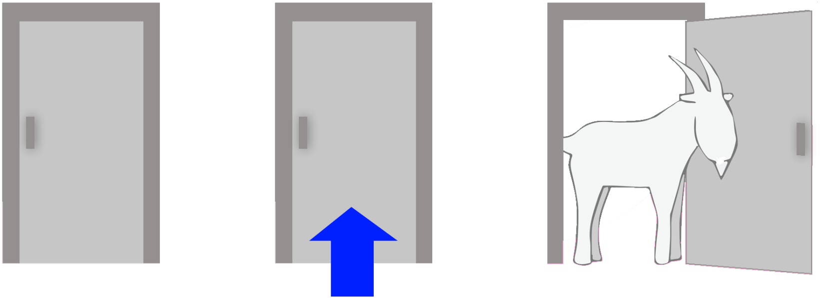

Initial probability:

1/3	1/3	1/3

New probability:

?	?	0

Formalizing the problem

Step 1 Outcome X is randomly drawn from $\mathcal{X} = \{x_1, x_2, x_3\}$ (the three doors) according to the uniform distribution p

Formalizing the problem

Step 1 Outcome X is randomly drawn from $\mathcal{X} = \{x_1, x_2, x_3\}$ (the three doors) according to the uniform distribution p

Step 2 The quizmaster, knowing X, chooses a set $Y \in \mathcal{Y} = \{\{x_1, x_2\}, \{x_2, x_3\}\}$ such that $Y \ni X$

- The structure of $\mathcal{Y}$ reflects that the quizmaster will always open one door, but never the door the contestant picked
- The chosen set Y is called the message

Formalizing the problem

Step 1 Outcome X is randomly drawn from $\mathcal{X} = \{x_1, x_2, x_3\}$ (the three doors) according to the uniform distribution p

Step 2 The quizmaster, knowing X, chooses a set $Y \in \mathcal{Y} = \{\{x_1, x_2\}, \{x_2, x_3\}\}$ such that $Y \ni X$

- The structure of $\mathcal{Y}$ reflects that the quizmaster will always open one door, but never the door the contestant picked
- The chosen set Y is called the message

Step 3 The contestant sees Y but not X, and responds with a prediction Q: his estimate of the currect probability distribution over $\mathcal{X}$

Generalizing the problem

We also want to know what probabilities to assign to the outcomes in a more general situation:

- For arbitrary (but finite) outcome spaces $\mathcal{X}$;
- For arbitrary marginal distribution p;
- For arbitrary families of allowed messages $\mathcal{Y}$.

The quizmaster's freedom of choice

- The quizmaster may use randomness when deciding which message Y to give us
- However, we don't know what distribution $P(Y \mid X)$ he uses

The quizmaster's freedom of choice

- The quizmaster may use randomness when deciding which message Y to give us
- However, we don't know what distribution $P(Y \mid X)$ he uses
- The conditional distribution P together with the marginal distribution p on $\mathcal{X}$ gives a joint distribution:

Quizmaster uses fair coin:

P	x_1	x_2	x_3
$\{x_1, x_2\}$	1/3	1/6	–
$\{x_2, x_3\}$	–	1/6	1/3
p_x	1/3	1/3	1/3

Quizmaster always opens x_3:

P	x_1	x_2	x_3
$\{x_1, x_2\}$	1/3	1/3	–
$\{x_2, x_3\}$	–	0	1/3
p_x	1/3	1/3	1/3

Worst-case approach

- We don't want to assume anything about how the quizmaster decides what message to give us
  - In the case of the original Monty Hall game, the fair-coin assumption can be defended based on the symmetry of the problem, but there may not be any symmetry in the general case
- Worst-case approach: we want to give guarantees on our predictions that hold no matter what mechanism is used to choose the message
- Corresponds to a two-player zero-sum game between the contestant and the quizmaster, where the quizmaster tries to make the contestant's prediction task as hard as possible

Loss functions

- To do this, we need some way of quantifying how good a prediction is
- We use a loss function that maps actual outcome x and contestant's prediction Q to some loss value for the contestant
- Examples:

$$\begin{aligned} \text{Logarithmic loss:} \quad & L(x, Q) = -\log Q(x) \\ \text{Brier loss:} \quad & L(x, Q) = \sum_{x' \in \mathcal{X}} (Q(x') - \mathbf{1}_{x'=x} \\ \text{Randomized 0–1 loss:} \quad & L(x, Q) = 1 - Q(x) \end{aligned}$$

Nash equilibrium

- The value of the game is

$$\sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} P(x, y) L(x, Q_y)$$

- The contestant chooses $(Q_y)_{y \in \mathcal{Y}}$ to minimize this; the quizmaster chooses P to maximize this
- For many loss functions, this game has a Nash equilibrium: there exists a pair of strategies for both players such that neither player would improve their situation by changing their strategy

Optimality theorems

- We proved optimality conditions for a very general class of loss functions
- For the special case that L is logarithmic loss, the characterization of optimality takes a very nice form:

Theorem

For logarithmic loss, a joint distribution P^* is optimal for the quizmaster if and only if there exists a vector $q \in [0,1]^{\mathcal{X}}$ such that

$$q_x = P^*(x \mid y) \text{ for all } x \in y \in \mathcal{Y} \text{ with } P^*(y) > 0, \text{ and}$$

$$\sum_{x \in y} q_x \leq 1 \text{ for all } y \in \mathcal{Y}$$

- We call this condition on P^* the RCAR condition

Optimal strategy may depend on the loss function

P	x_1	x_2	x_3	x_4
$\{x_1, x_2\}$	$1/3$	$1/6$	—	—
$\{x_2, x_3\}$	—	$1/6$	$1/6$	$1/6$
p_x	$1/3$	$1/3$	$1/6$	$1/6$

- This strategy P is optimal for logarithmic loss (it satisfies the RCAR condition), but not for Brier loss

RCAR condition beyond log loss

- If the set of available messages $\mathcal{Y}$ forms a graph (meaning that each message contains exactly two outcomes), then the RCAR condition characterizes optimality regardless of the loss function;
- If $\mathcal{Y}$ forms a matroid (satisfies the matroid basis exchange property), then the same is true;
- For any other $\mathcal{Y}$, this is not the case: there exists some marginal p such that the optimal strategies for log loss and Brier loss are different

Conclusion

- Worst-case optimal probability updating gives us a general framework that we can use to update our beliefs about an outcome in situations where the joint distribution of outcome and evidence is unknown
- In general, the answer depends on the loss function ("subjective"), but in certain situations they don't ("objective")