Updating Probabilities –
a new look at some paradoxes

Peter Grünwald

Centrum Wiskunde & Informatica – Amsterdam

Mathematisch Instituut – Universiteit Leiden

partially based on joint work with
J. Halpern (2003, Journal of Artificial Intelligence Research)
and R. Gill (2008, Annals of Statistics)

most recent contribution: G., ECSQARU 2013

Is conditioning always appropriate?

- How should probabilities be updated in the light of new information?
- Standard answer: conditioning!
- But…

  ‘naïve’ conditioning sometimes gives ‘wrong’ results

    - Monty Hall (quizmaster, 3-door, car & goat)
    - 3-prisoners puzzle
    - 2-children puzzle

Standard Updating: Conditioning

1. I throw fair die; I see outcome $X \in \{1, ..., 6\}$ but you don't
2. I tell you whether $X \in \{1, 2, 3, 4\}$ or not.
3. You have to come up with distribution for X (e.g., to predict X) How do you do this?

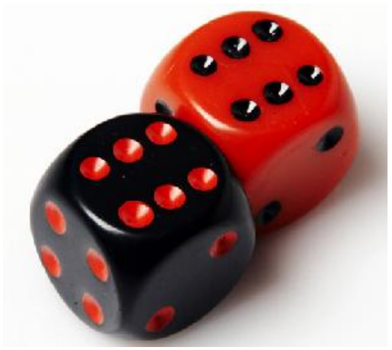

Standard Updating: Conditioning

1. I throw fair die; I see outcome $X \in \{1, \ldots, 6\}$ but you don't
2. I tell you whether $X \in \{1, 2, 3, 4\}$ or not.
3. You have to come up with distribution for X (e.g., to predict X) How do you do this?
   - you condition on $\{1, 2, 3, 4\}$ so that, e.g.

$$P(X = 4 \mid X \in \{1, \ldots, 4\}) = \frac{P(X = 4, X \in \{1, 2, 3, 4\})}{P(X \in \{1, 2, 3, 4\})} =$$

$$\frac{P\{4\}}{P\{1, 2, 3, 4\}} = \frac{\frac{1}{6}}{\frac{4}{6}} = \frac{1}{4}.$$

The Three Prisoners Problem

Gardner 1959

- There are three prisoners, **A**, **B** en **C**
- Two of them will be chosen at random and will be executed. The prisoners know this.
- **A** will be executed with probability 2/3, so he survives with probability 1/3.

The Three Prisoners Problem

- Rita, the jailor, walks by. A asks her: can you tell me whether B or C will be executed?
- Rita says: **B**
- It seems that Rita does not give A any new information about his survival probability
  - A already knew that either B or C would be executed

The Three Prisoners Problem

- Rita, the jailor, walks by. A asks her: can you tell me whether B or C will be executed?
- Rita says: **B**
- It seems that Rita does not give A any new information about his survival probability
  - A already knew that either B or C would be executed
  - With this reasoning, the probability that A survives remains 1/3

- But: if you condition, probability becomes 1/2 !
- First there were three possibilities with equal probability:

  A survives, **B** survives, **C** survives.

- After jailer says “B will be executed”, only two possibilities left:

  A survives, **C** survives.

- So the probability that **A** survives is now 1/2 !

- But: if you condition, probability becomes 1/2 !
- First there were three possibilities with equal probability:

  A survives, **B** survives, **C** survives.
- After jailer says “B will be executed”, only two possibilities left:

  A survives, **C** survives.
- So the probability that **A** survives is now 1/2 !

What's wrong?

1. If we condition, the probability that A survives increases, no matter what the jailer answers!

What's wrong?

1. If we condition, the probability that A survives increases, no matter what the jailer answers!
   - That can't be right!
2. Alternative reasoning: the probability that A survives remains equal, no matter what the jailer answers!
   - That's not entirely correct either…

What’s wrong?

1. If we condition, the probability that A survives increases, no matter what the jailer answers!
   – That can’t be right!
2. Alternative reasoning: the probability that A survives remains equal, no matter what the jailer answers!
   – That’s not entirely correct either…

- The best answer is:
  you cannot say any more what the probability is
  – Unless you make extra assumptions about the jailor’s mind…(then 2 may be correct; 1 not)

Main Insight

- **Something's wrong with "naïve" conditioning**
  - 3 Prisoners, Monty Hall, 2- Children…

Main Insight:
Naïve vs. Sophisticated Spaces

- **Something's wrong with "naïve" conditioning**
  - 3 Prisoners, Monty Hall, 2- Children…
- The real problem is not conditioning, but representing the situation by an overly simple, 'naïve' space
- Conditioning *is* always appropriate…in the 'sophisticated' space

Main Insight:
Naïve vs. Sophisticated Spaces

- **Something's wrong with "naïve" conditioning**
  - 3 Prisoners, Monty Hall, 2- Children…
- The real problem is not conditioning, but representing the situation by an overly simple, 'naïve' space
- Conditioning *is* always appropriate…in the 'sophisticated' space
- Sophisticated space takes protocol into account; naïve space doesn't
  - What would the jailer do in the one case that he has a choice? (if A is not executed)

Example Nr 0.

1. I throw fair die; I see outcome $X \in \{1, \ldots, 6\}$ but you don't
2. I tell you whether $X \in \{1, 2, 3, 4\}$ or not.
3. You have to come up with distribution for X (e.g., to predict X) How do you do this?

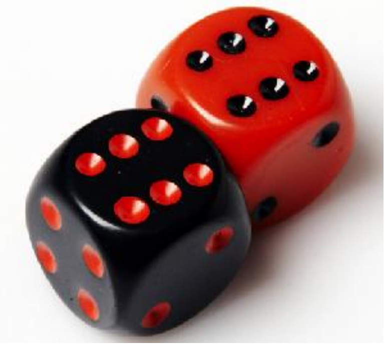

Example nr. 1

1. I throw fair die; I see outcome $X \in \{1, \ldots, 6\}$ but you don't
2. I tell you either $X \in \{1, 2, 3, 4\}$ or $X \in \{3, 4, 5, 6\}$
3. You have to come up with distribution for X

Example nr. 1

1. I throw fair die; I see outcome $X \in \{1, \ldots, 6\}$ but you don't
2. I tell you either $X \in \{1, 2, 3, 4\}$ or $X \in \{3, 4, 5, 6\}$
3. You have to come up with distribution for X

Now conditioning can be very misleading!

$$P(X = 4 \mid X \in \{1, \ldots, 4\}) = \frac{1}{4}$$

…but I may decide that whenever $X = 4$

I tell you "$X \in \{3, 4, 5, 6\}$"

Example nr. 1

- Whenever you "observe" $X \in \{1, 2, 3, 4\}$ you predict
$$P(X = 4 \mid X \in \{1, 2, 3, 4\}) = \frac{1}{4}$$
but in fact, it may be that $X = 4$ never occurs in such a case
  - depending on my protocol, the 'true' probability that $X = 4$ can be anything between 0 and 1/3
  - to get a precise probability, you must know my protocol!

Naive vs. Sophisticated Conditioning

- Naive Conditioning
  - first put probability distribution on set of outcomes (e.g. $\{1, \ldots, 6\}$)
  - then condition relative to that set

Naive vs. Sophisticated Conditioning

- Naive Conditioning
  - first put probability distribution on set of outcomes (e.g. $\{1, \ldots, 6\}$)
  - then condition relative to that set
- Sophisticated Conditioning
  - first put probability distribution on ***larger* set** which includes protocol
    (e.g. 2 additional events
    - 'if I observe 4, I say $\{1, \ldots, 4\}$' p
    - 'if I observe 4, I say $\{3, \ldots, 6\}$' $1 - p$
  - then condition relative to that larger space

3 Prisoner's Revisited

- Naive Conditioning
  - first put probability distribution on set of outcomes (e.g. $\{A, B, C\}$)
  - then condition relative to that set
- Sophisticated Conditioning
  - first put probability distribution on ***larger* set** which includes protocol
    (e.g. 2 additional events
    - 'if A survives, jailer says $\{A, C\}$, i.e. "B dies" ' p
    - 'if A survives, jailer says $\{A, B\}$, i.e. "C dies" ' $1 - p$
  - then condition relative to that larger space

Naive vs. Sophisticated

- Naive Conditioning
  - **what people tend to do...**
    ...but usually incorrect! (G&H 2003)

- Sophisticated Conditioning
  - works fine...
    ...but often cannot be applied because protocol may be unknown

Our Work

- When do we get the correct results by ignoring the protocol, i.e. working in the naive space?

Our Work

- When do we get the correct results by ignoring the protocol, i.e. working in the naive space?
- Answer: a *sufficient* condition is:
  if the set of possible observations has no **overlap!**

$\{1,2,3,4\},\{3,4,5,6\}$ ⟶ **bad!**

$\{1,3,5\},\{2,4,6\}$ ⟶ **good!**

Our Work

- When do we get the correct results by ignoring the protocol, i.e. working in the naive space?
- Answer: a *sufficient* condition is:
  if the set of possible observations has no **overlap!**

Main Insight: Many probabilistic paradoxes arise because people tend to represent problem in a ‘naive’ space, in which the set of observations has overlap. ***Then conditioning makes no sense!***

Monty Hall (3-door) Problem

Monty Hall 1970

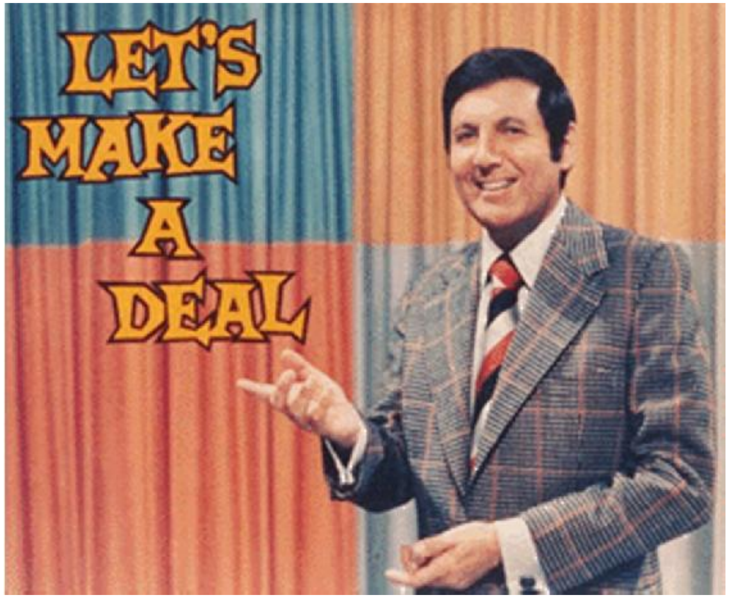

Monty Hall (3-door) Problem

Monty Hall 1970

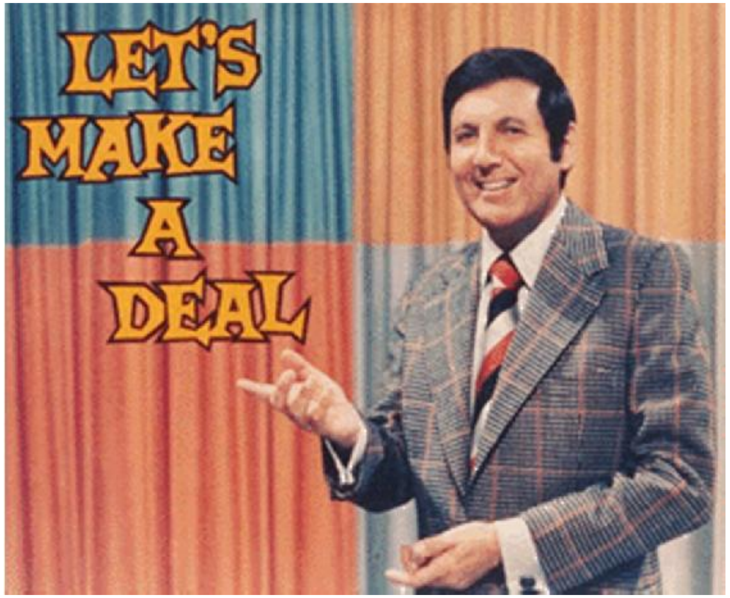

in the Nederlands also known as
"**Willem Ruis Probleem**"

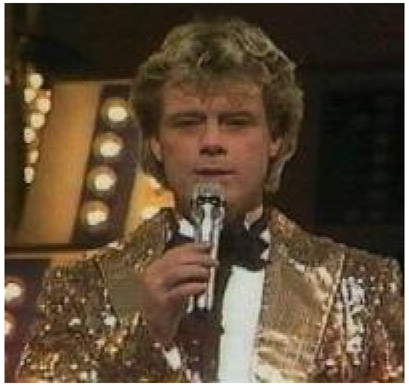

Monty Hall

- There are three doors in the TV studio. Behind one door is a car, behind both other doors a goat. You choose one of the doors. Monty Hall opens one of the other two doors, and shows that there is a goat behind it. You are now allowed to switch to the other door that is still closed. Is it smart to switch?

Monty Hall

- Switching doors is very smart
- Yet almost everybody initially thought that it doesn´t matter whether you switch or not, since they feel that for both doors that are still closed, the probability that the car is behind it, is ½.

Monty Hall

- Switching doors is very smart
- Yet almost everybody initially thought that it doesn´t matter whether you switch or not, since they feel that for both doors that are still closed, the probability that the car is behind it, is ½.
- Most treatments of this puzzle explain why the answer ½ is wrong
- I am interested in a different question:
  why does everyone make this mistake?
  Answer: because it is the result of naïve conditioning!

Monty Hall

- Most treatments of this puzzle explain why the answer ½ is wrong
- I am interested in a different question: why does everyone make this mistake?
  Answer: because it is the result of naïve conditioning!
  outcome space: $\{A, B, C\}$
  possible observations: $\{A, B\}$, $\{A, C\}$
  people reason:

$$P(A) = P(B) = P(C) = \frac{1}{3}$$

$$P(A \mid \text{door C opened}) = P(A \mid \{A, B\}) = \frac{1}{2}$$

Monty Hall

- Most treatments of this puzzle explain why the answer ½ is wrong
- I am interested in a different question: why does everyone make this mistake?

  Answer: because it is the result of naïve conditioning!

  outcome space: $\{A, B, C\}$

  possible observations: $\{\mathbf{A}, B\}$, $\{\mathbf{A}, C\}$

$$P(A \mid \text{door C opened}) \overset{?}{=} P(A \mid \{A, B\}) = \frac{1}{2}$$

There is overlap in set of observations so we can immediately infer that this is highly suspicious!

Second Course:
Two Children Puzzle

- Mrs. Jones has two children. One of them is a boy. What’s the probability that the other child is a boy?

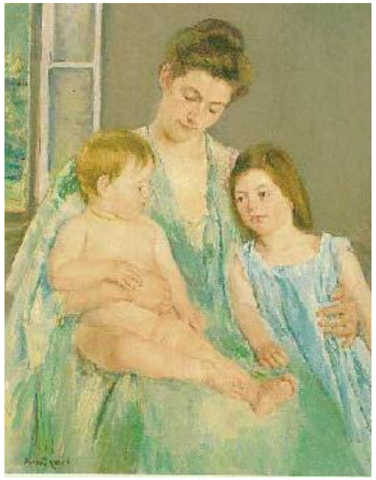

Second Course:
Two Children Puzzle

- Mrs. Jones has two children. One of them is a boy. What's the probability that the other child is a boy?
- Correct answer is: **1/3** (and not ½ !)

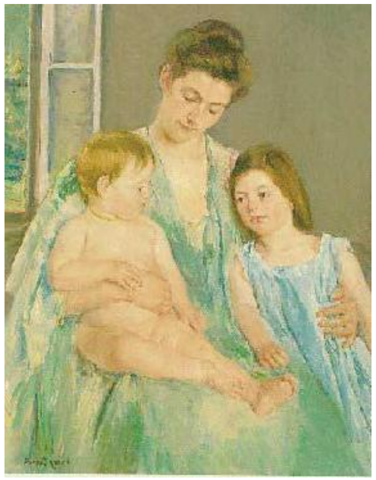

Two Children Puzzle

- Mrs. Jones has two children. One of them is a boy. What's the probability that the other child is a boy?
- "Correct" answer is: **1/3**

$$P(\{(B,B)\} \mid \{(B,G),(G,B),(B,B)\}) = \frac{1}{3}$$

Two Children Puzzle

- Mrs. Jones has two children. One of them is a boy. What's the probability that the other child is a boy?
- "Correct" answer is: **1/3**

$$P(\{(B,B)\} \mid \{(B,G),(G,B),(B,B)\}) = \frac{1}{3}$$

OR IS IT?

Precise two Children Puzzle

- Suppose I pick a person P uniformly at random from population of parents with two children and I start conversation with P
- Scenario 1: I ask whether at least one of the children is a boy. Suppose answer is *yes.* Then the probability of 2 boys is indeed 1/3

$$\textbf{yes} \Leftrightarrow \{(B,G),(G,B),(B,B))\}$$

$$\textbf{no} \Leftrightarrow \{(G,G))\}$$

$$P(\{(B,B)\} \mid \textbf{yes}) = P(\{(B,B)\} \mid \{(B,G),(G,B),(B,B)\}) = \frac{1}{3}$$

Precise two Children Puzzle

- Suppose I pick a person P uniformly at random from population of parents with two children and I start conversation with P
- Scenario 1: I ask whether at least one of the children is a boy. Suppose answer is *yes.* Then the probability of 2 boys is indeed 1/3

$$\textbf{yes} \Leftrightarrow \{(B,G),(G,B),(B,B))\}$$

$$\textbf{no} \Leftrightarrow \{(G,G))\}$$

- Scenario 2: P spontaneously says something that implies that she has at least one boy (like “my Johnny just got his zwemdiploma”). Then 1/3 *not* correct

$$\textbf{mentions boy} \Leftrightarrow \{(B,G),(G,B),(B,B)\}$$

$$\textbf{mentions girl} \Leftrightarrow \{(B,G),(G,B),(G,G)\}$$

Overlap!

Two Children Puzzle 2.0

- I pick person P uniformly at random from population of parents with two children and I start conversation with P
- I ask whether at least one of the children is a boy. Suppose answer is *yes.* The probability of 2 boys is now 1/3 (uncontroversially)
- Suppose I ask instead: is your youngest child a boy? Suppose answer is *yes.* Now the probability of 2 boys is 1/2 (uncontroversially)

Two Children Puzzle 2.0

- I pick person P uniformly at random from population of parents with two children and I start conversation with P
- I ask whether at least one of the children is a boy. Suppose answer is *yes.* The probability of 2 boys is now 1/3 (uncontroversially)
- Suppose I ask instead: is your youngest child a boy? Suppose answer is *yes.* Now the probability of 2 boys is 1/2 (uncontroversially)
- Suppose I ask instead: is your oldest child a boy? Suppose answer is *yes.* Now the probability of 2 boys is 1/2 (uncontroversially)

Two Children Puzzle 2.0

- I pick person P uniformly at random from population of parents with two children and I start conversation with P
- I ask whether at least one of the children is a boy. Suppose answer is *yes.* The probability of 2 boys is now 1/3 (uncontroversially)
- I ask: can you give me the name of your son (if you have two sons, just give me one name, you can choose which!)
- Suppose P says: "yes: Martin" (prob. of two boys is still 1/3)
- I ask: is Martin your *youngest* child?
  - Answer: "yes".

  Now you know youngest child is boy, so $\Pr(2 \text{ Boys}) = 1/2$

  - Answer: "no".

  Now you know oldest child is a boy, so $\Pr(2 \text{ Boys}) = 1/2$

$$\frac{1}{2} = \frac{1}{3} \text{????}$$

Our Earlier Work

- When do we get the correct results by ignoring the protocol, i.e. working in the naive space?
- Answer: a *sufficient* condition is:
  if the set of possible observations has no **overlap!**

Main Insight: Many probabilistic paradoxes arise because people tend to represent problem in a ‘naive’ space, in which the set of observations has overlap. ***Then conditioning makes no sense!***

Our Earlier Work

- When do we get the correct results by ignoring the protocol, i.e. working in the naive space?
- Answer: a *sufficient* condition is:
  if the set of possible observations has no **overlap!**

This condition is essentially also a necessary condition, except for some very special cases (G&H ’03, G&G ’08)

Our Earlier Work

- When do we get the correct results by ignoring the protocol, i.e. working in the naive space?
- Answer: a *sufficient* condition is:
  if the set of possible observations has no **overlap!**

First time the importance of this idea was stressed was by Glenn Shafer (1985), who proved that, if the set of messages you can be told is a partition, then conditioning is the right thing to do in a Dutch-book-like argument (conditioning as a theorem rather than a definition!)

A (New) Radical Proposal

- Standard Notation for Event-Based Conditioning is
  $$P(A \mid B)$$
  where B can be any measurable set
- I think we should change probability theory!
  change not(at)ion of event-based conditioning to
  $$P(A \mid B \| \mathcal{P})$$
  which is only defined if $B \in \mathcal{P}$ and $\mathcal{P}$ is a **partition** of the sample space

A (New) Radical Proposal

- Standard Notation for Event-Based Conditioning is

  $$P(A \mid B)$$

  where B can be any measurable set

- I think we should change probability theory!

  change not(at)ion of event-based conditioning to

  $$P(A \mid B \| \mathcal{P})$$

  which is only defined if $B \in \mathcal{P}$ and $\mathcal{P}$ is a **partition** of the sample space

Note that probabilists **usually** don't have to worry about this because they condition on random variables and then $\mathcal{P}$ is implicitly given! (you know that the set of events on which you could possibly conditions forms a partition anyway)

A (New) Radical Proposal

- But this is only the beginning...
- ...see Grünwald (2013, ECSQARU) for a proposal what to do if you **cannot** condition in the standard sense!
  - Semi-Naive, strongly Bayesian Answer: act as if Monty throws a fair coin to decide what door to open, if he has a choice
  - Standard but unsatisfactory Non-Naive Answer: **dilation**
  - There is third way!

Monty Hall – The Wikipedia Wars

- I disagree with the common opinion that "if you switch, the probability of winning is 2/3".
- The actual probability of winning depends on Monty's protocol, which we don´t know.
- Only if Monty flips a fair coin to determine whether to open B or C if he has a choice, is the probability really 2/3.
- Most people (including Vos Savant) implicitly assume this in their analysis of Monty Hall…but why!?!?

Monty Hall – The Wikipedia Wars

- A better analysis (according to me, Richard Gill, and some of the wikipedia warriors) is to make no assumptions about Monty's behaviour.
- Then the probability of winning if you switch can be anything between ½ and 1. Thus, we do agree that switching is better than staying! (it's a dominating strategy)

Dilation

- For standard Monty Hall game (specify whether to switch or not), second analysis is convincing
- But more generally, it is still problematic:
  - Before Monty opens a door, you assess the probability that you win if you stay at A as 1/3
  - If Monty opens door B, then you would update this probability to be 'inbetween 0 and ½'
  - If Monty opens door C, you would also update this probability to be 'inbetween 0 and ½'
  - No matter what you observe, you have less knowledge about A after you observe it than before. Can this be right!?!?

dilation, Seidenfeld and Wasserman, 1993

A (New) Radical Proposal

- But this is only the beginning...
- ...see Grünwald (2013, ECSQARU) for a proposal what to do if you **cannot** condition in the standard sense!
  - Semi-Naive, strongly Bayesian Answer: act as if Monty throws a fair coin to decide what door to open, if he has a choice
  - Standard but unsatisfactory Non-Naive Answer: **dilation**
  - There is a third way!

Conclusion

- To determine whether conditioning on event A is a meaningful operation, we need to know what other things than ‘A’ we might have been told.
  - also applies to “two-envelope paradox”
- Unfortunately, you will not find this observation in any probability text book!
- …ideas also have applications in statistics
  - missing data, survival analysis
- …and there should also be relation to **belief updating in logic** (see B. Kooij, 2003, J. van Benthem, 2009)