Randomness and imprecision

Gert de Cooman   Jasper De Bock

Ghent University, SYSTeMS

gert.decooman@UGent.be
http://users.UGent.be/~gdcooma
gertekoo.wordpress.com

UTOPIAE TS-II Public Lecture
4 July 2018

UTOPIAE
Uncertainty Treatment and Optimisation in Aerospace Engineering
Handling the unknown at the edge of tomorrow

http://utopiae.eu
http://twitter.com/utopiae_network
info@utopiae.eu

1 1 0 0 1 0 1 0 1 1 0 1 0 1 0 1 0 0 0 1 1 1 0 1 0 1 0 1 1

1 1 1 0 0 1 1 1 0 0 1 0 0 1 1 1 0 0 1 1 1 1 0 1 0 1 0 1 1 1 . . .

WHEN IS A SEQUENCE RANDOM?

Random sequences and random numbers

$$1100101011010101000111010101 1$$

$$111001110010011100111101010111\ldots$$

Random sequences and random numbers

$$
\begin{aligned}
0.&11001010110101010001110101011\\
&111001110010011100111101010111\ldots
\end{aligned}
$$

is a real number in $[0,1]$.

Random sequences and calibrated forecasts

0 1 1 0 0 1 0 1 0 . . .

Random sequences and calibrated forecasts

$$\tfrac{1}{2}\,0 \quad \tfrac{1}{2}\,1 \quad \tfrac{1}{2}\,1 \quad \tfrac{1}{2}\,0 \quad \tfrac{1}{2}\,0 \quad \tfrac{1}{2}\,1 \quad \tfrac{1}{2}\,0 \quad \tfrac{1}{2}\,1 \quad \tfrac{1}{2}\,0 \ldots$$

Random sequences and calibrated forecasts

p_1 0 p_2 1 p_3 1 p_4 0 p_5 0 p_6 1 p_7 0 p_8 1 p_9 0 ...

Random sequences and calibrated forecasts

$I_1$0 $I_2$1 $I_3$1 $I_4$0 $I_5$0 $I_6$1 $I_7$0 $I_8$1 $I_9$0 . . .

A BIT OF HISTORY

The classical case of a fair coin

$$\tfrac{1}{2}0 \quad \tfrac{1}{2}1 \quad \tfrac{1}{2}1 \quad \tfrac{1}{2}0 \quad \tfrac{1}{2}0 \quad \tfrac{1}{2}1 \quad \tfrac{1}{2}0 \quad \tfrac{1}{2}1 \quad \tfrac{1}{2}0 \ldots$$

A bit of notation

$$\omega = (x_1, x_2, x_3, \dots, x_n, \dots) \in \Omega$$

with $\Omega = \{0,1\}^{\mathbb{N}} \approx [0,1]$

$$\omega^n = (x_1, x_2, x_3, \dots, x_n) \in \Omega^{\Diamond}$$

with $\Omega^{\Diamond} = \{0,1\}^*$

$$\omega_n = x_n \in \{0,1\}$$

THE APPROACH OF VON MISES, WALD AND CHURCH

The approach of von Mises, Wald and Church

Randomness of ω means:

$$\frac{\sum_{k=1}^{n} x_k}{n} \to \frac{1}{2} \qquad \text{(Law of Large Numbers)}$$

The approach of von Mises, Wald and Church

Randomness of ω means:

$$\frac{\sum_{k=1}^{n} x_k}{n} \to \frac{1}{2} \qquad \text{(Law of Large Numbers)}$$

but also more stringently, for any selection rule $S\colon \{0,1\}^* \to \{0,1\}$ in a countable class $\mathscr{S}$:

$$\frac{\sum_{k=1}^{n} S(x_1,\ldots,x_{k-1})x_k}{\sum_{k=1}^{n} S(x_1,\ldots,x_{k-1})} \to \frac{1}{2}$$

$$\text{whenever } \sum_{k=1}^{n} S(x_1,\ldots,x_{k-1}) \to \infty$$

A selection rule S is a way of selecting subsequences from ω:

$$\begin{cases} S(x_1,\ldots,x_{k-1}) = 1 & \Rightarrow \text{select } x_k \\ S(x_1,\ldots,x_{k-1}) = 0 & \Rightarrow \text{discard } x_k \end{cases}$$

The approach of von Mises, Wald and Church

For von Mises and Wald, $\mathscr{S}$ represented the countable class of selection rules that can be constructed in some given formal system of arithmetic.

For Church, $\mathscr{S}$ represented the countable class of computable selection rules.

$\Rightarrow$ Computable stochasticity

On both approaches, there is an uncountable infinity of ‘random’ sequences ω associated with $\mathscr{S}$: they have (Lebesgue) measure one on $[0, 1]$.

Criticism

Jean Ville in his *Étude critique de la notion de collectif* (1939):

There are other limit laws than the Law of Large Numbers that are not implied by Computable Stochasticity,

e.g. oscillation around the limit.

Computable stochasticity seems too weak!

THE MARTIN-LÖF APPROACH

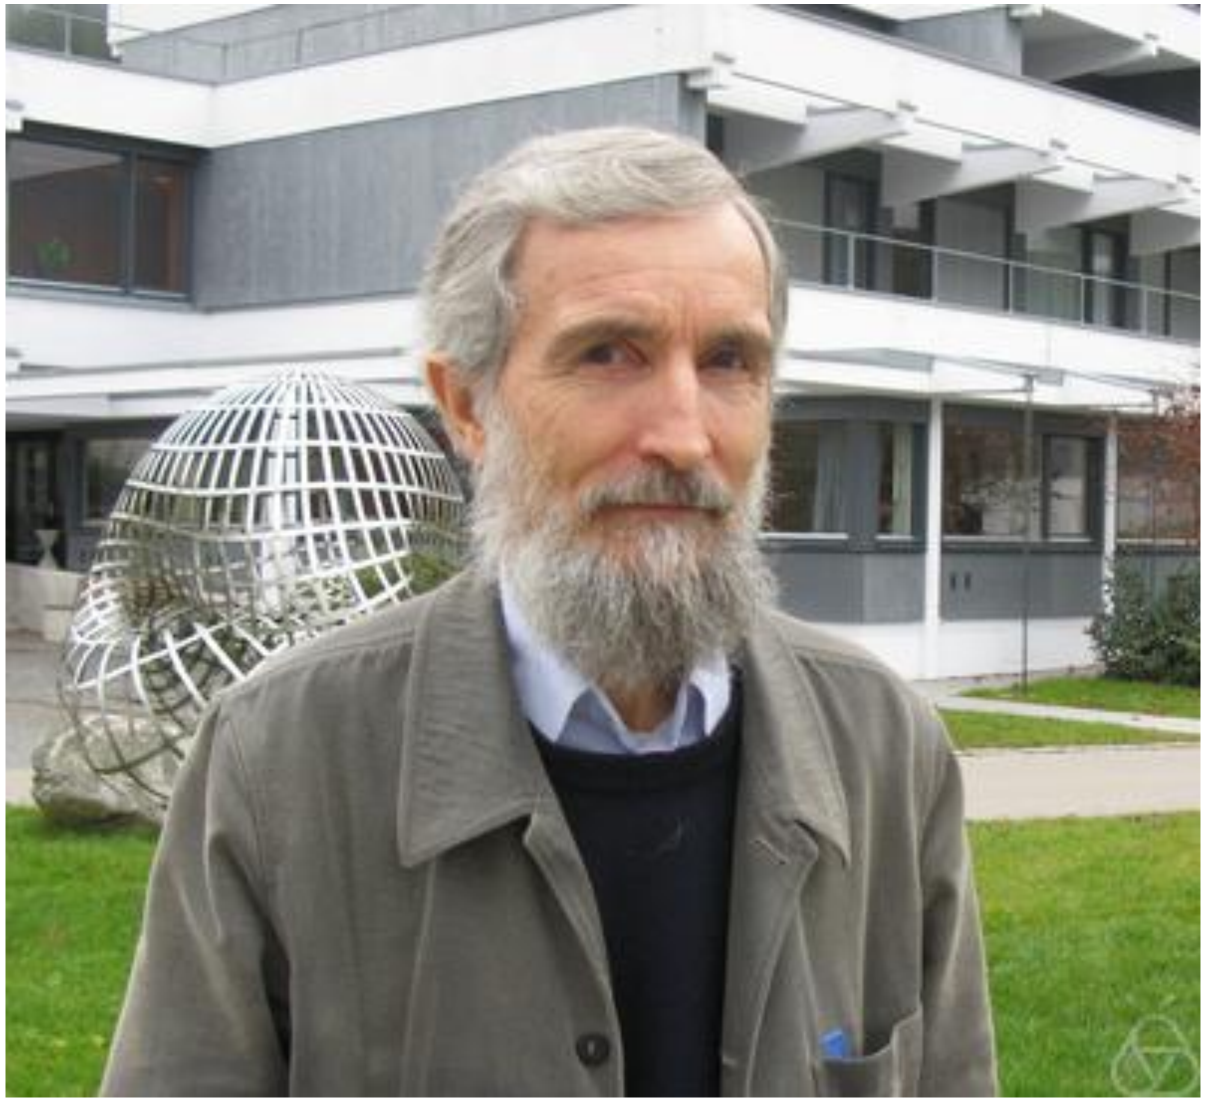

Martin-Löf randomness and avoiding null sets

Basic observations:

- randomness is about satisfying limit laws
- randomness is therefore about avoiding null sets
- only countably many null sets can be avoided
- only countably many can be constructed
- a subset A of $[0,1]$ is null if for all $\varepsilon > 0$ there is a sequence of intervals covering A with total measure at most ε

Effectively null set

A subset A of $[0,1]$ is effectively null if there is an algorithm that turns any rational $\varepsilon > 0$ into a sequence of intervals covering A with total measure at most ε.

Martin-Löf randomness and avoiding null sets

Conclusions:

- there are only countably many effectively null sets
- their union is null, so its complement has measure one.

Martin-Löf randomness

A sequence ω is Martin-Löf random if it belongs to no effectively null set.

The Martin-Löf random sequences have measure one,
and they are computably stochastic.

FORECASTING AND THE MARTINGALE APPROACH

More general precise forecasting

p_1 0 p_2 1 p_3 1 p_4 0 p_5 0 p_6 1 p_7 0 p_8 1 p_9 0 . . .

A single precise forecast r

Forecaster
specifies his expectation r for an unknown outcome X in $\{0, 1\}$:
his commitment to adopt r as a fair price for X.

A single precise forecast r

Forecaster

specifies his expectation r for an unknown outcome X in $\{0, 1\}$:
his commitment to adopt r as a fair price for X.

Skeptic

takes Forecaster up on his commitments:

(i) for any $p \leq r$ and $\alpha \geq 0$, Forecaster must accept $\alpha(X - p)$;

(ii) for any $q \geq r$ and $\beta \geq 0$, Forecaster must accept $\beta(q - X)$.

A single precise forecast r

Forecaster

specifies his expectation r for an unknown outcome X in $\{0, 1\}$:
his commitment to adopt r as a fair price for X.

Skeptic

takes Forecaster up on his commitments:

(i) for any $p \leq r$ and $\alpha \geq 0$, Forecaster must accept $\alpha(X - p)$;

(ii) for any $q \geq r$ and $\beta \geq 0$, Forecaster must accept $\beta(q - X)$.

Reality

determines the value x of X.

Gambles available to Skeptic: precise forecast r

$$f(X) = -\alpha(X-p) - \beta(q-X) \text{ with } \alpha, \beta \geq 0 \text{ and } 0 \leq p \leq r \leq q \leq 1$$

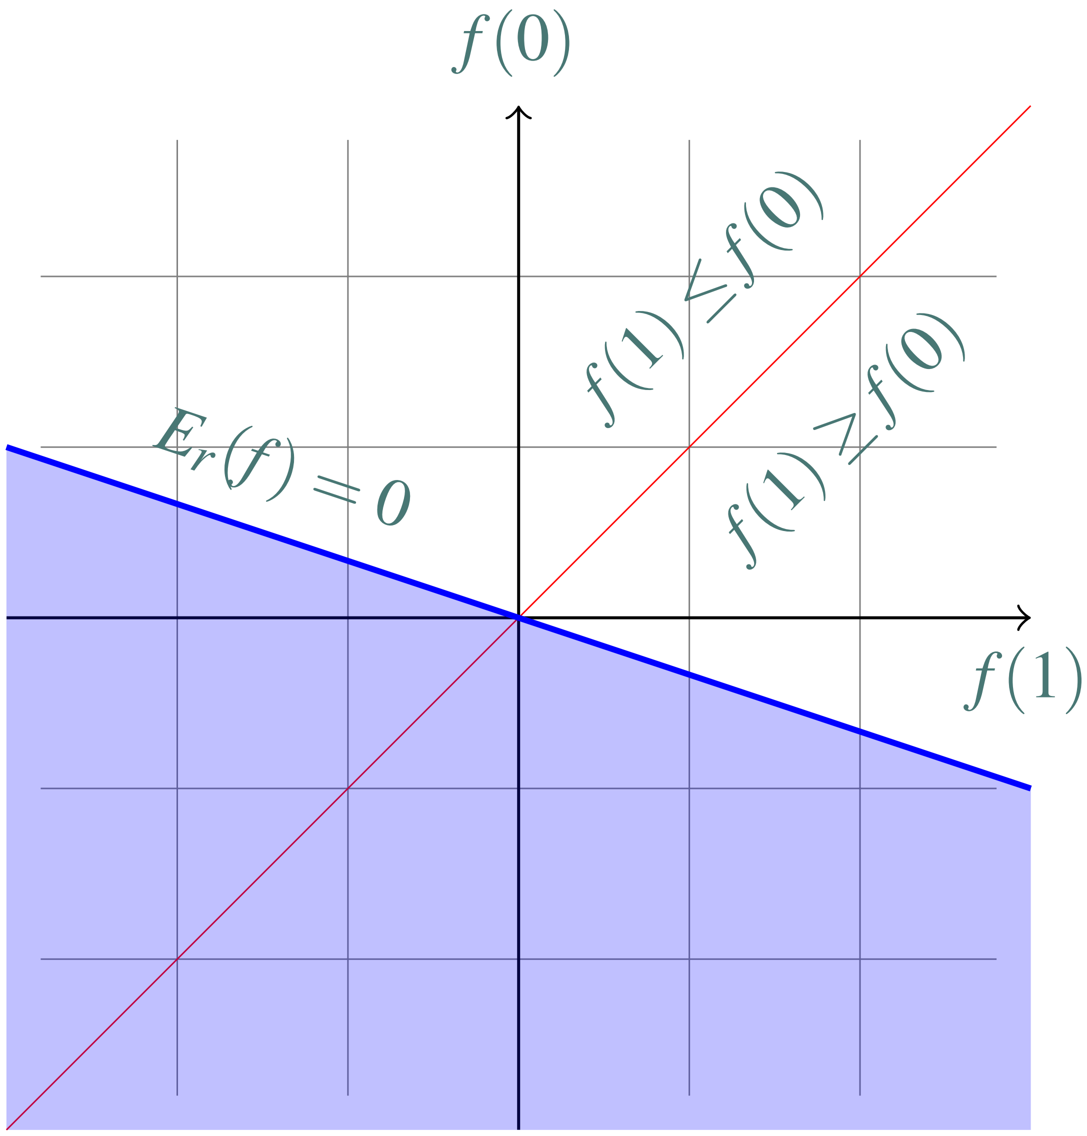

$$E_r(f) := rf(1) + (1-r)f(0) \leq 0$$

More forecasts: event tree

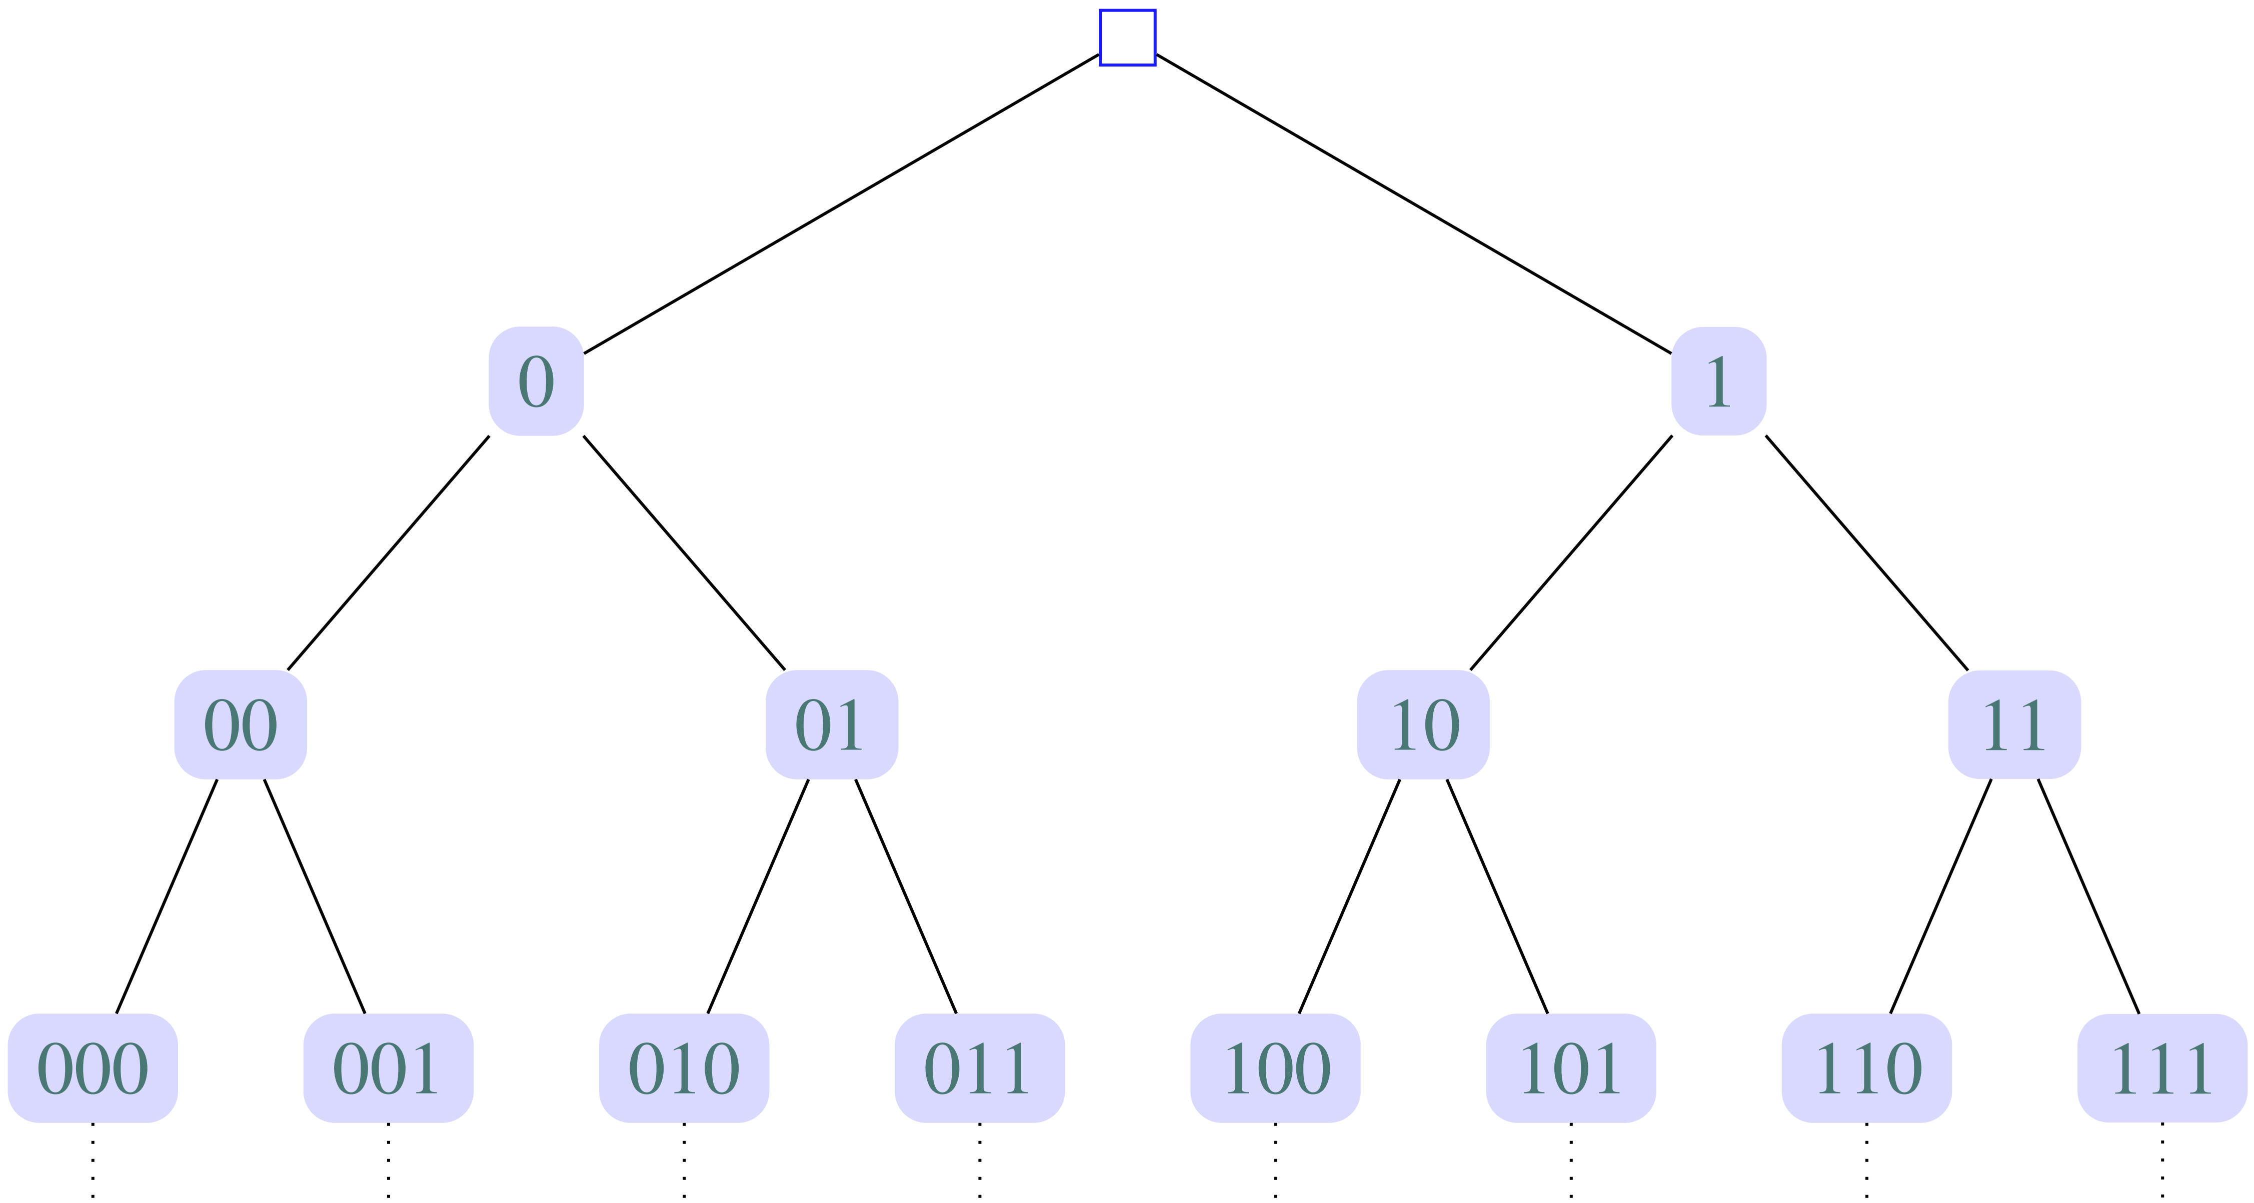

More forecasts: probability tree

In a probability tree, we associate a precise forecast $\gamma(s) = p_s$ with each situation $s \in \Omega^\lozenge$:

forecasting system $\gamma\colon \Omega^\lozenge \to [0,1]$

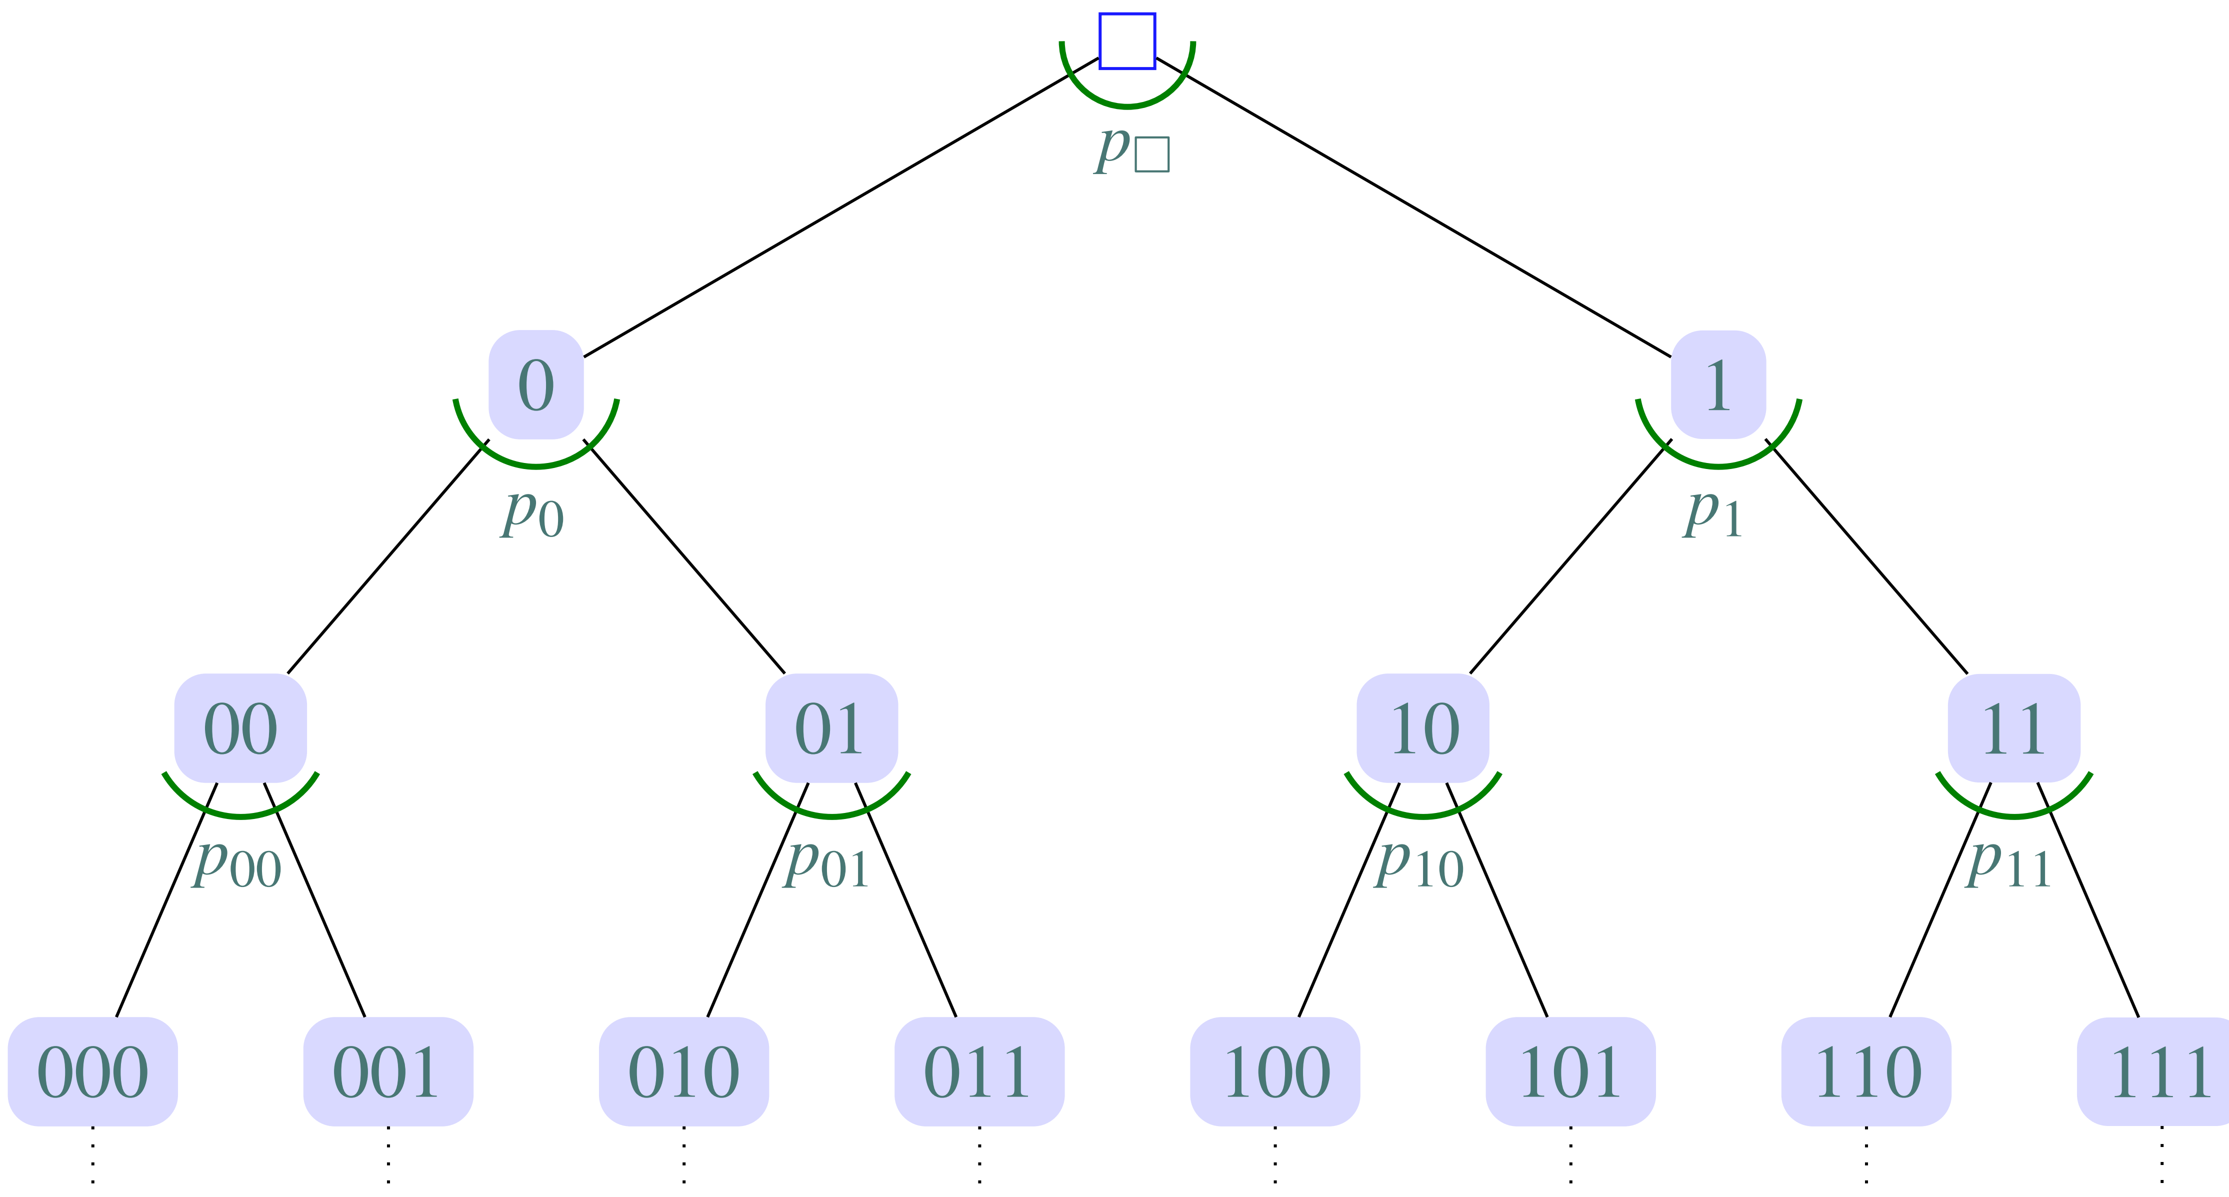

Event trees and processes

A real process is a map $M\colon \Omega^{\Diamond} \to \mathbb{R}$, so attaches a real number $M(s)$ to every situation s.

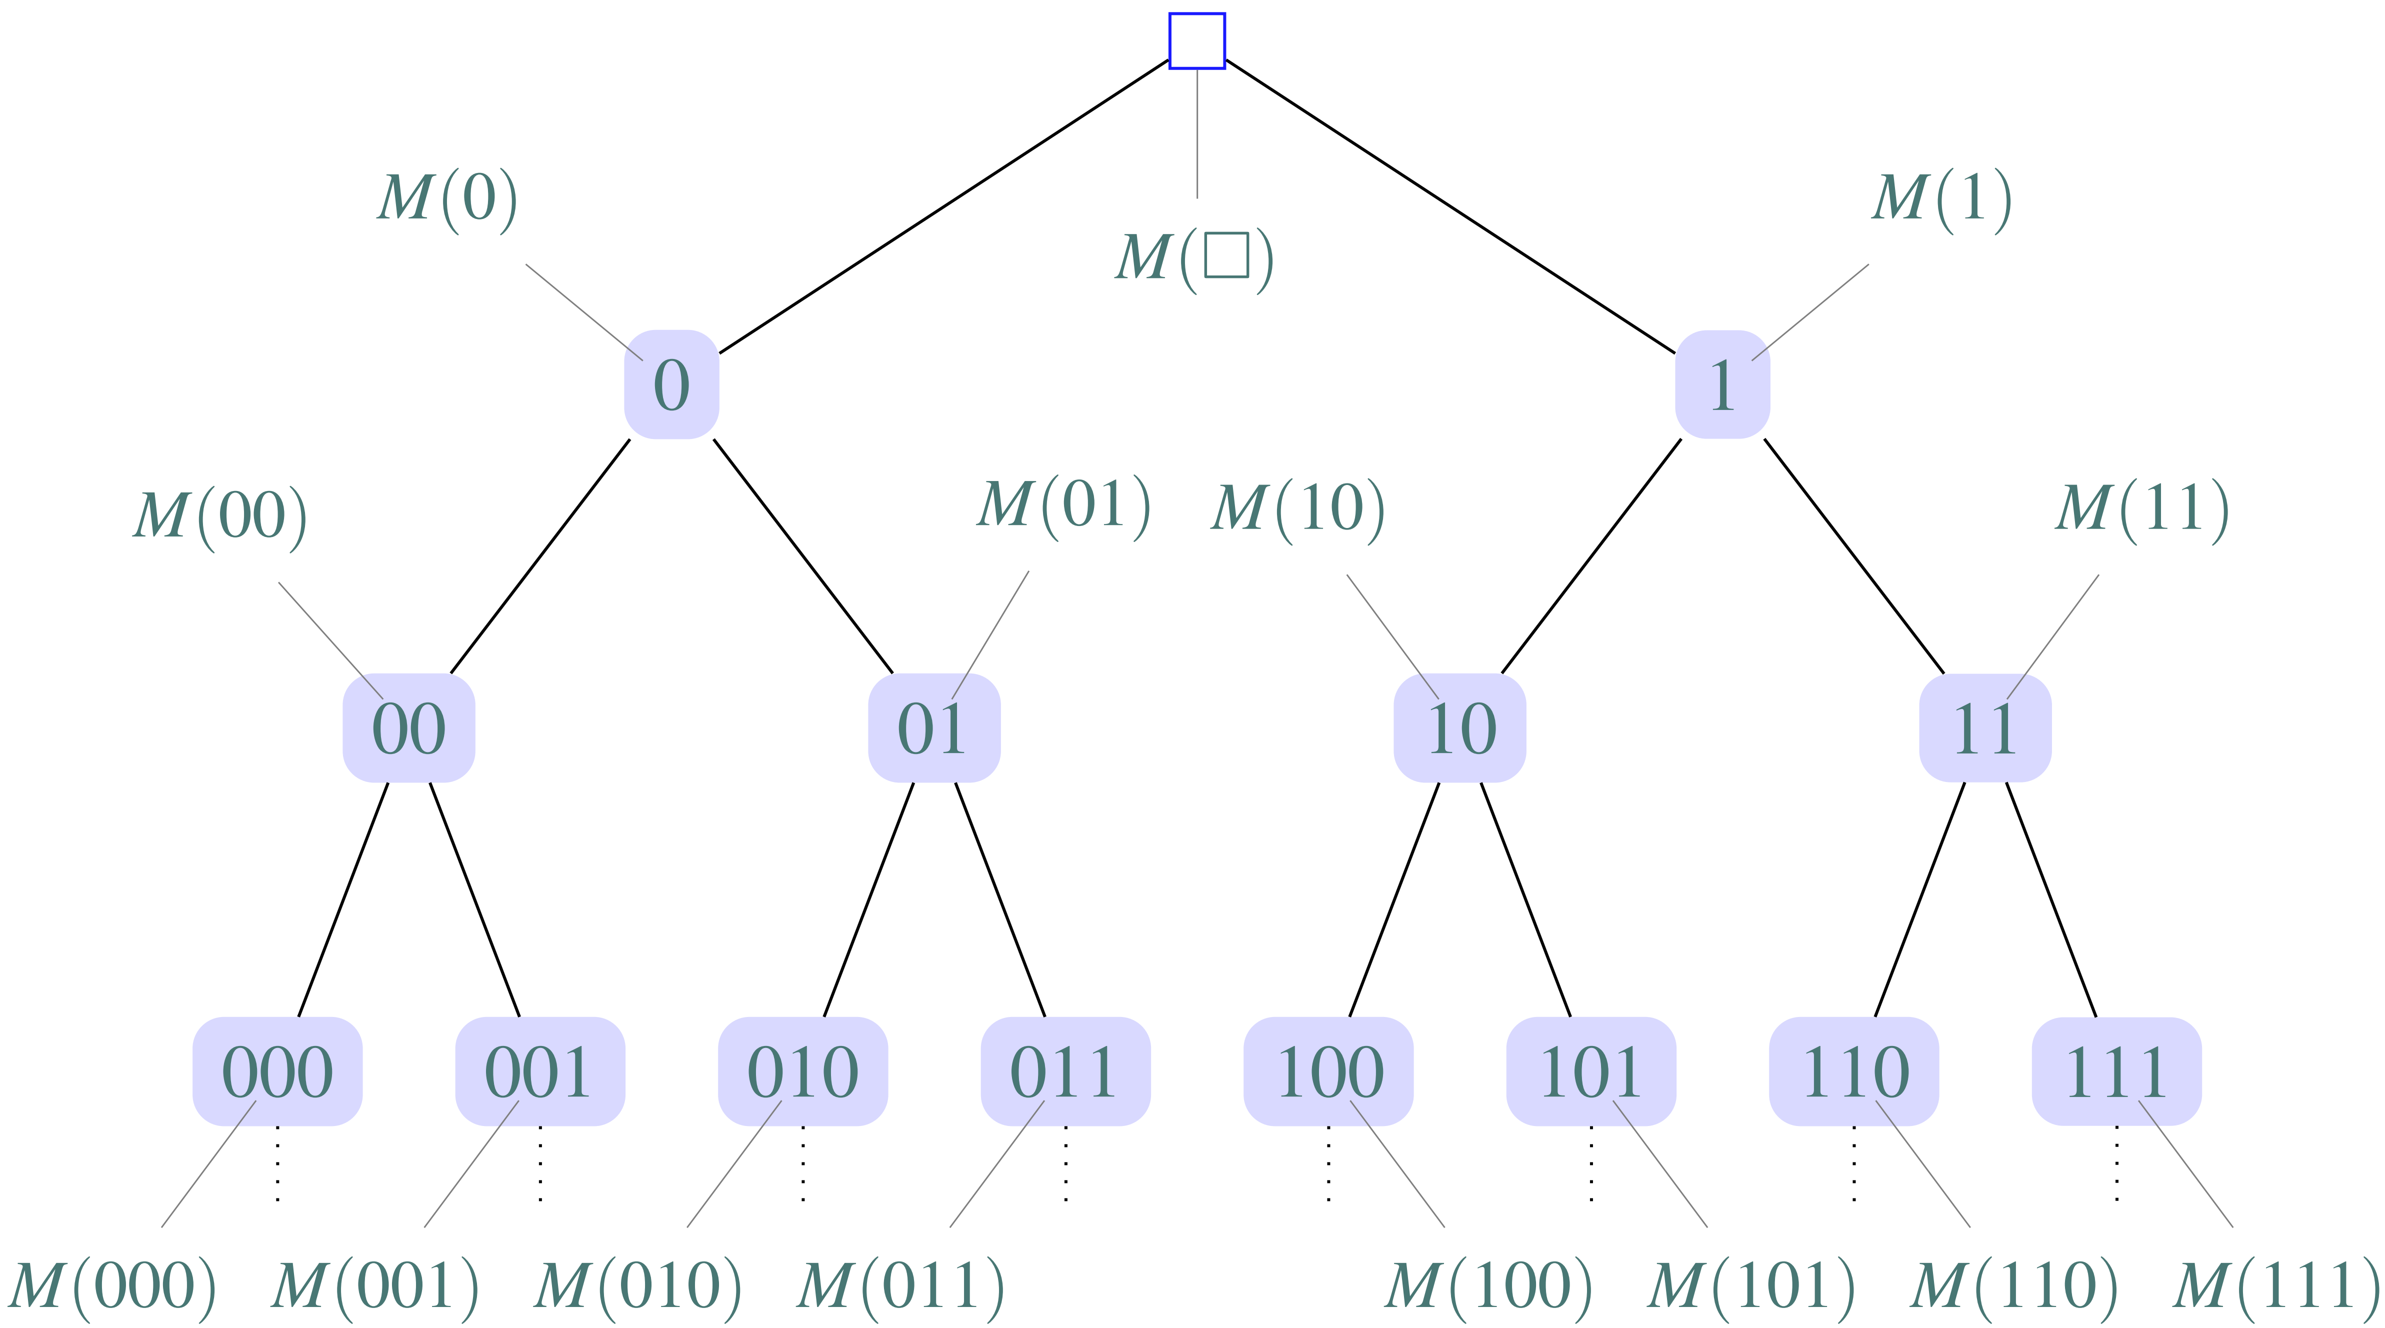

Probability tree and supermartingales

A capital process M for Skeptic is the result of his taking up an available gamble f_s in every possible situation s:

$$\left.\begin{aligned} M(s1) &= M(s) + f_s(1) \\ M(s0) &= M(s) + f_s(0) \end{aligned}\right\} \text{ with } E_s(f_s) \leq 0$$

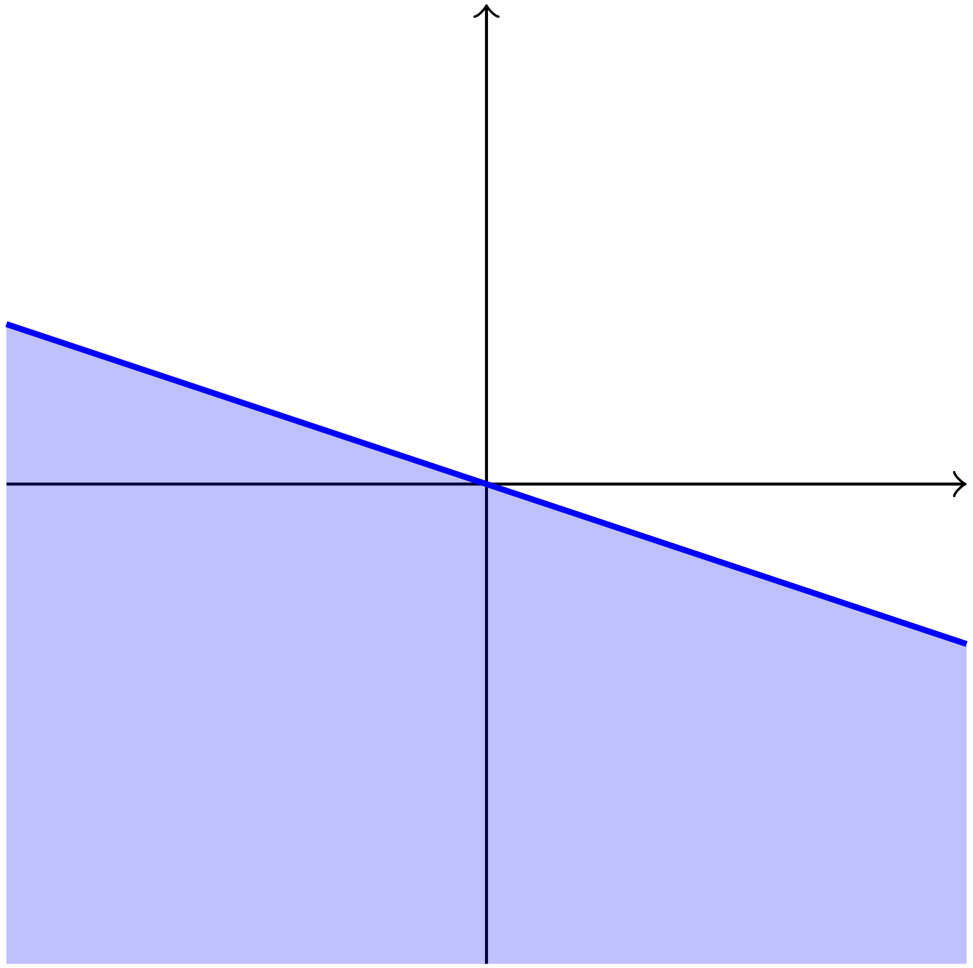

Probability tree and supermartingales

A capital process M for Skeptic is the result of his taking up an available gamble f_s in every possible situation s:

$$\left.\begin{aligned} M(s1) &= M(s) + f_s(1) \\ M(s0) &= M(s) + f_s(0) \end{aligned}\right\} \text{ with } E_s(f_s) \leq 0$$

Supermartingale

A supermartingale M for a forecasting system γ is a real process whose increments

$$\Delta M(s) := M(s\cdot) - M(s)$$

have non-positive expectation:

$$E_{\gamma(s)}(\Delta M(s)) \leq 0 \text{ in all situations } s.$$

Probability tree, supermartingales and expectations

Jean Ville's Theorem (1939)

For any measurable bounded function $g : [0,1] \to \mathbb{R}$:

$$E_\gamma(g) = \inf\left\{ M(\square) : M \text{ supermartingale and } \liminf_{n\to+\infty} M(\omega^n) \geq g(\omega) \right\}$$

Randomness on the martingale approach

The essential idea idea behind randomness is that there is no system for breaking the bank, for becoming unboundedly rich by betting on the successive outcomes in the sequence.

Randomness

A sequence ω is random for a forecasting system γ if no *non-negative* allowable supermartingale for γ becomes unbounded on ω.

Randomness on the martingale approach

The essential idea idea behind randomness is that there is no system for breaking the bank, for becoming unboundedly rich by betting on the successive outcomes in the sequence.

Martin-Löf randomness

A sequence ω is Martin-Löf random for a forecasting system γ if no *non-negative* lower semicomputable supermartingale for γ becomes unbounded on ω.

Randomness on the martingale approach

The essential idea idea behind randomness is that there is no system for breaking the bank, for becoming unboundedly rich by betting on the successive outcomes in the sequence.

Computable randomness

A sequence ω is computably random for a forecasting system γ if no *non-negative* computable supermartingale for γ becomes unbounded on ω.

Randomness on the martingale approach

The essential idea idea behind randomness is that there is no system for breaking the bank, for becoming unboundedly rich by betting on the successive outcomes in the sequence.

Schnorr randomness

A sequence ω is Schnorr random for a forecasting system γ if no *non-negative* computable supermartingale for γ becomes computably unbounded on ω.

ALLOWING FOR IMPRECISION

More general precise forecasting

$I_1$0 $I_2$1 $I_3$1 $I_4$0 $I_5$0 $I_6$1 $I_7$0 $I_8$1 $I_9$0 . . .

A single interval forecast $I = [\underline{p}, \overline{p}]$

Forecaster
specifies his interval forecast $I = [\underline{p}, \overline{p}]$ for an unknown outcome X in $\{0, 1\}$: his commitment to adopt $\underline{p}$ as a highest buying price and $\overline{p}$ as a lowest selling price for X.

A single interval forecast $I = [\underline{p}, \overline{p}]$

Forecaster

specifies his interval forecast $I = [\underline{p}, \overline{p}]$ for an unknown outcome X in $\{0,1\}$: his commitment to adopt $\underline{p}$ as a highest buying price and $\overline{p}$ as a lowest selling price for X.

Skeptic

takes Forecaster up on his commitments:

(i) for any $p \leq \underline{p}$ and $\alpha \geq 0$, Forecaster must accept $\alpha(X - p)$;

(ii) for any $q \geq \overline{p}$ and $\beta \geq 0$, Forecaster must accept $\beta(q - X)$.

A single interval forecast $I = [\underline{p}, \overline{p}]$

Forecaster

specifies his interval forecast $I = [\underline{p}, \overline{p}]$ for an unknown outcome X in $\{0, 1\}$: his commitment to adopt $\underline{p}$ as a highest buying price and $\overline{p}$ as a lowest selling price for X.

Skeptic

takes Forecaster up on his commitments:

(i) for any $p \leq \underline{p}$ and $\alpha \geq 0$, Forecaster must accept $\alpha(X - p)$;

(ii) for any $q \geq \overline{p}$ and $\beta \geq 0$, Forecaster must accept $\beta(q - X)$.

Reality

determines the value x of X.

Gambles available to Skeptic: interval forecast

$$I = [\underline{p}, \overline{p}]$$

$f(X) = -\alpha(X - p) - \beta(q - X)$ with $\alpha, \beta \geq 0$ and $0 \leq p \leq \underline{p} \leq \overline{p} \leq q \leq 1$

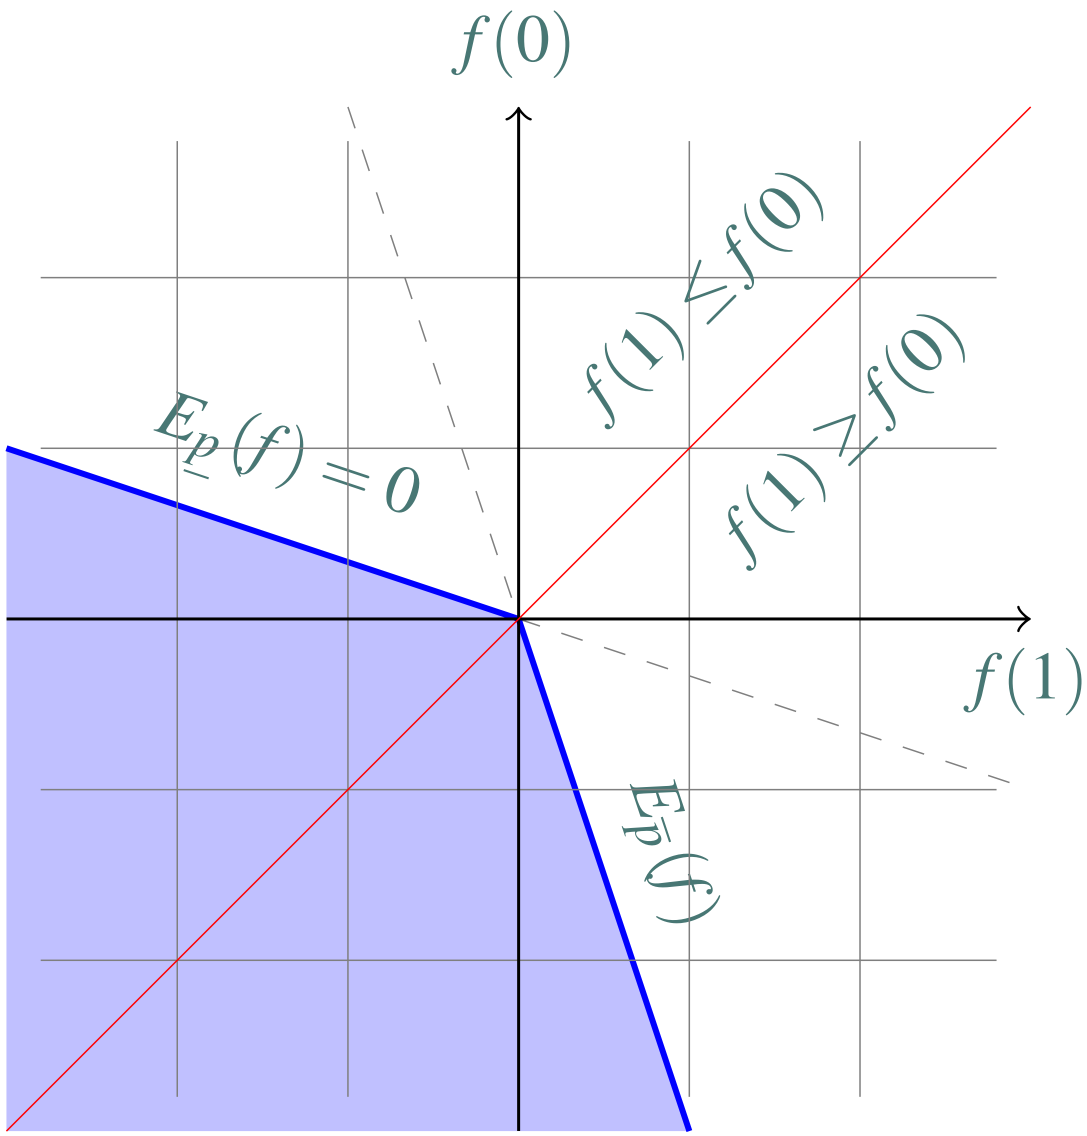

$$\overline{E}_I(f) := \max_{r \in I} E_r(f) \leq 0$$

Gambles available to Skeptic: vacuous forecast

$$I = [0,1]$$

$f(X) = -\alpha(X-p) - \beta(q-X)$ with $\alpha, \beta \geq 0$ and $0 = p$ and $q = 1$

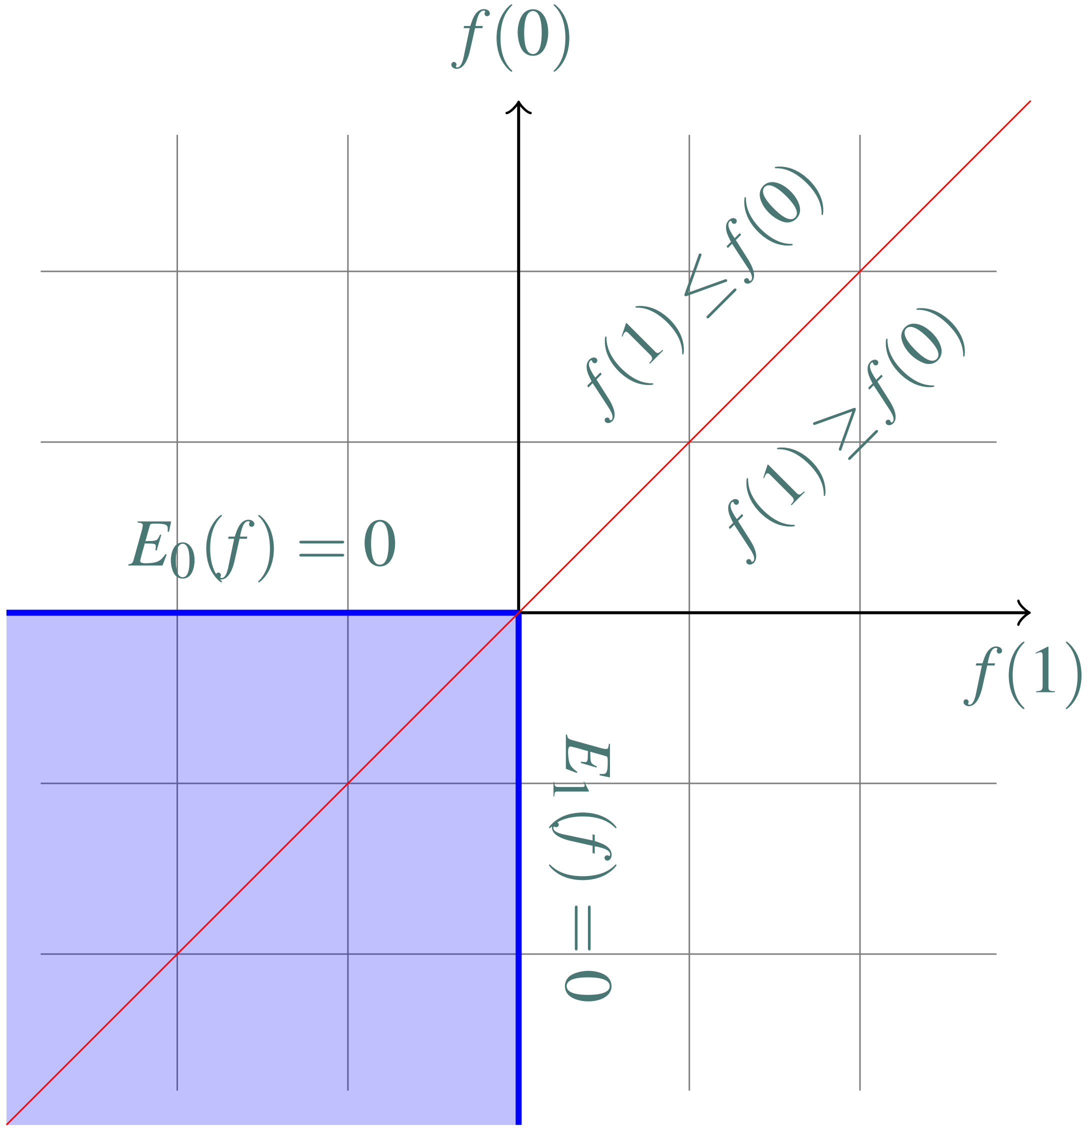

$$\overline{E}_I(f) := \max_{r \in [0,1]} E_r(f) = \max f \leq 0$$

More forecasts: imprecise probability tree

In an imprecise probability tree, we associate an interval forecast $\gamma(s) = I_s = [\underline{p}_s, \overline{p}_s]$ with each situation $s \in \Omega^\Diamond$:

forecasting system $\gamma\colon \Omega^\Diamond \to \mathscr{C}$

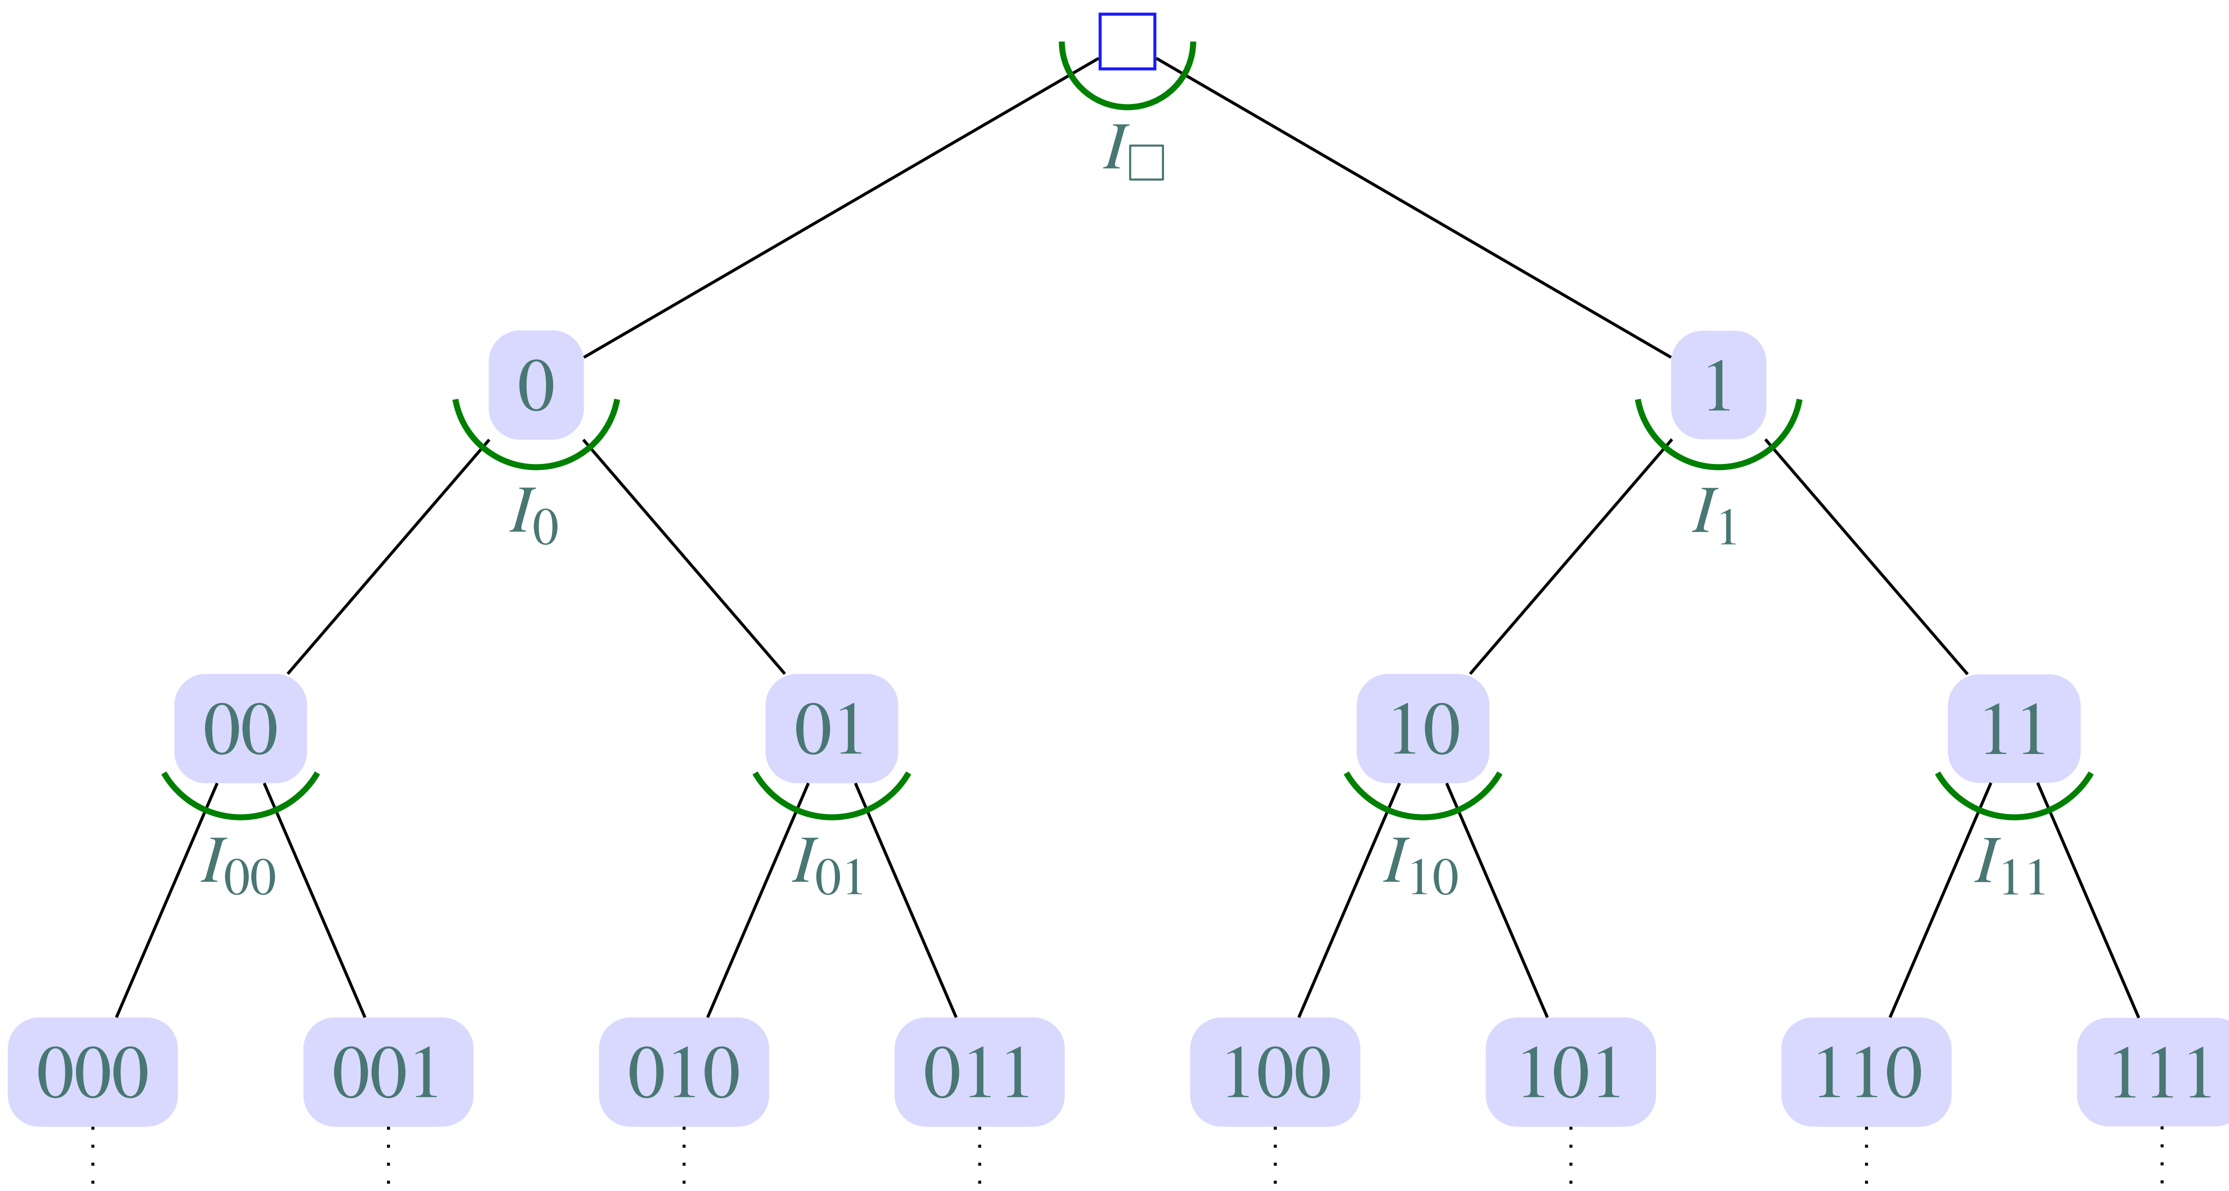

Imprecise probability tree and supermartingales

A capital process M for Skeptic is the result of his taking up an available gamble f_s in every possible situation s:

$$\left.\begin{aligned} M(s1) &= M(s) + f_s(1) \\ M(s0) &= M(s) + f_s(0) \end{aligned}\right\} \text{ with } \overline{E}_s(f_s) \leq 0$$

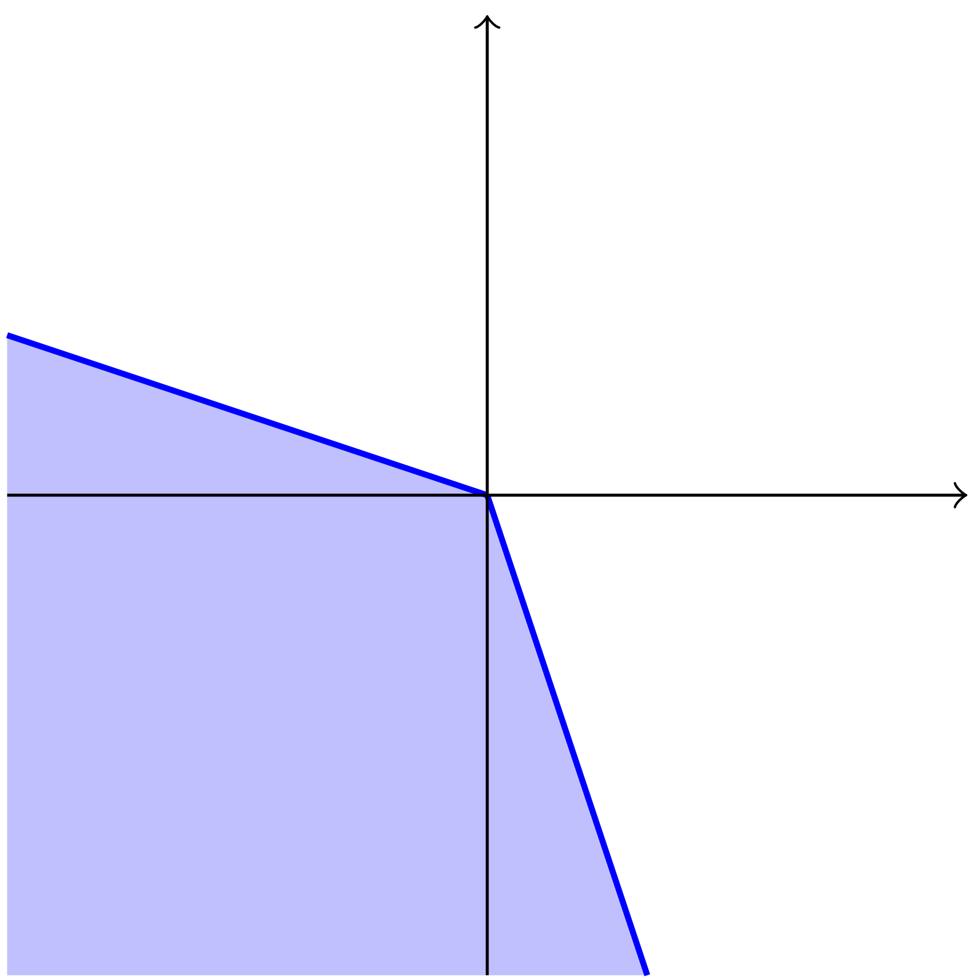

Imprecise probability tree and supermartingales

A capital process M for Skeptic is the result of his taking up an available gamble f_s in every possible situation s:

$$\left.\begin{aligned} M(s1) &= M(s) + f_s(1) \\ M(s0) &= M(s) + f_s(0) \end{aligned}\right\} \text{ with } \overline{E}_s(f_s) \leq 0$$

Supermartingale

A supermartingale M for a forecasting system γ is a real process whose increments

$$\Delta M(s) := M(s\cdot) - M(s)$$

have non-positive upper expectation:

$$\overline{E}_{\gamma(s)}(\Delta M(s)) \leq 0 \text{ in all situations } s.$$

Imprecise probability tree, supermartingales and upper expectations

Jean Ville's Theorem (1939)

For a precise forecasting system γ, and for any measurable bounded function $g\colon [0,1] \to \mathbb{R}$:

$$E_\gamma(g) = \inf\Big\{M(\square) : M \text{ supermartingale and } \liminf_{n\to+\infty} M(\omega^n) \geq g(\omega)\Big\}$$

Upper expectation

For an imprecise forecasting system γ, and any bounded function $g\colon [0,1] \to \mathbb{R}$:

$$\overline{E}_\gamma(g) = \inf\Big\{M(\square) : M \text{ supermartingale and } \liminf_{n\to+\infty} M(\omega^n) \geq g(\omega)\Big\}$$

Randomness on the martingale approach

The essential idea idea behind randomness is that there is no system for breaking the bank, for becoming unboundedly rich by betting on the successive outcomes in the sequence.

Randomness

A sequence ω is random for a forecasting system γ if no *non-negative* allowable supermartingale for γ becomes unbounded on ω.

Randomness on the martingale approach

The essential idea idea behind randomness is that there is no system for breaking the bank, for becoming unboundedly rich by betting on the successive outcomes in the sequence.

Martin-Löf randomness

A sequence ω is Martin-Löf random for a forecasting system γ if no *non-negative* lower semicomputable supermartingale for γ becomes unbounded on ω.

Randomness on the martingale approach

The essential idea idea behind randomness is that there is no system for breaking the bank, for becoming unboundedly rich by betting on the successive outcomes in the sequence.

Computable randomness

A sequence ω is computably random for a forecasting system γ if no *non-negative* computable supermartingale for γ becomes unbounded on ω.

Randomness on the martingale approach

The essential idea idea behind randomness is that there is no system for breaking the bank, for becoming unboundedly rich by betting on the successive outcomes in the sequence.

Schnorr randomness

A sequence ω is Schnorr random for a forecasting system γ if no *non-negative* computable supermartingale for γ becomes computably unbounded on ω.

CONSISTENCY RESULTS

Consistency

Every forecaster believes himself to be well-calibrated:

Consider any forecasting system $\gamma\colon \Omega^\Diamond \to \mathscr{C}$. Then (strictly) almost all outcome sequences are computably random for γ in the imprecise probability tree that corresponds to γ.

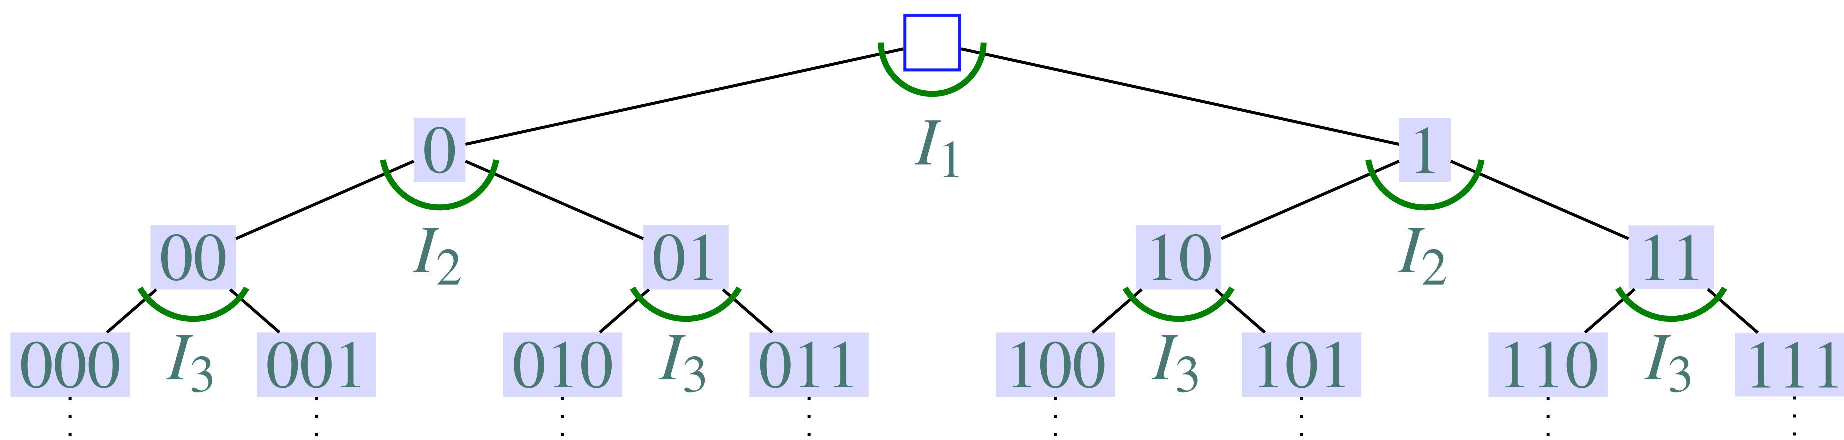

Corollary

For any sequence of interval forecasts $\phi = (I_1, \dots, I_n, \dots)$ there is a forecasting system such that (strictly) almost all outcome sequences are computably random for this forecasting system in the associated imprecise probability tree.

Constant interval forecasts

$$\gamma_I(s) := I \text{ for all } s \in \Omega^\lozenge.$$

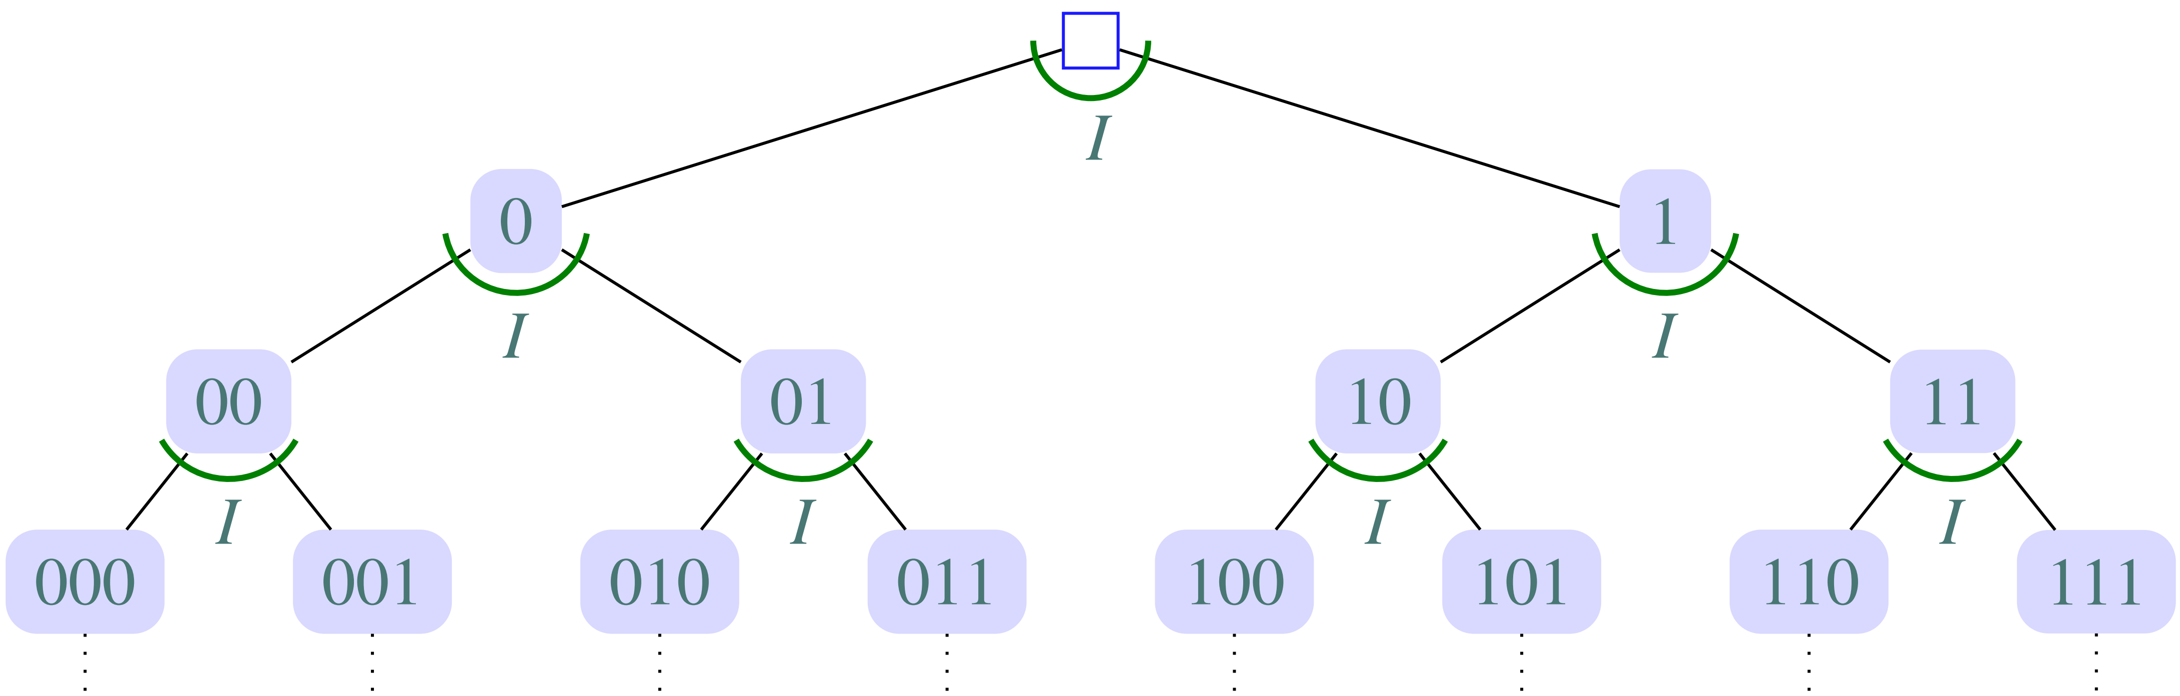

$$\mathscr{C}_{\mathrm{C}}(\omega) = \{I \in \mathscr{C} : \gamma_I \text{ makes } \omega \text{ computably random}\}$$

Church randomness or computable stochasticity

Theorem

Consider any outcome sequence $\omega = (x_1, \dots, x_n, \dots)$ in Ω and any constant interval forecast $I = [\underline{p}, \overline{p}] \in \mathscr{C}_{\mathbb{A}}(\omega)$ that makes ω (Martin-Löf or computably) random. Then for any computable selection rule $S \colon \Omega^{\Diamond} \to \{0, 1\}$ such that $\sum_{k=0}^{n} S(x_1, \dots, x_k) \to +\infty$:

$$\underline{p} \leq \liminf_{n \to +\infty} \frac{\sum_{k=0}^{n-1} S(x_1, \dots, x_k) x_{k+1}}{\sum_{k=0}^{n-1} S(x_1, \dots, x_k)}$$

$$\leq \limsup_{n \to +\infty} \frac{\sum_{k=0}^{n-1} S(x_1, \dots, x_k) x_{k+1}}{\sum_{k=0}^{n-1} S(x_1, \dots, x_k)} \leq \overline{p}.$$

RANDOMNESS IS INHERENTLY IMPRECISE

The set filter $\mathscr{C}_{\mathrm{C}}(\omega)$

Fix any ω in Ω.

Non-emptiness

$[0,1] \in \mathscr{C}_{\mathrm{C}}(\omega)$, so any sequence of outcomes ω has at least one stationary forecast that makes it computably random: $\mathscr{C}_{\mathrm{C}}(\omega) \neq \emptyset$.

The set filter $\mathscr{C}_{\mathrm{C}}(\omega)$

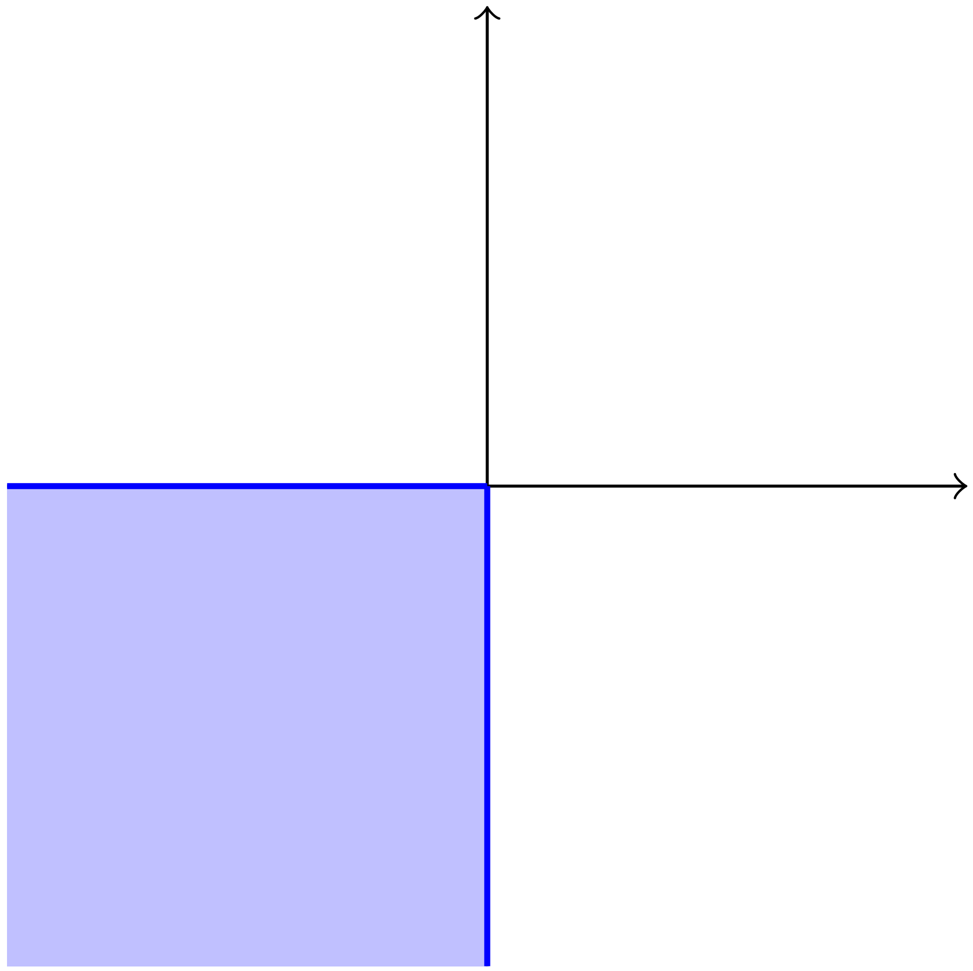

The set filter $\mathscr{C}_{\mathrm{C}}(\omega)$

Fix any ω in Ω.

Non-emptiness

$[0,1] \in \mathscr{C}_{\mathrm{C}}(\omega)$, so any sequence of outcomes ω has at least one stationary forecast that makes it computably random: $\mathscr{C}_{\mathrm{C}}(\omega) \neq \emptyset$.

Increasingness

For all $I, J \in \mathscr{C}$: if $I \in \mathscr{C}_{\mathrm{C}}(\omega)$ and $I \subseteq J$, then $J \in \mathscr{C}_{\mathrm{C}}(\omega)$.

The set filter $\mathscr{C}_{\mathrm{C}}(\omega)$

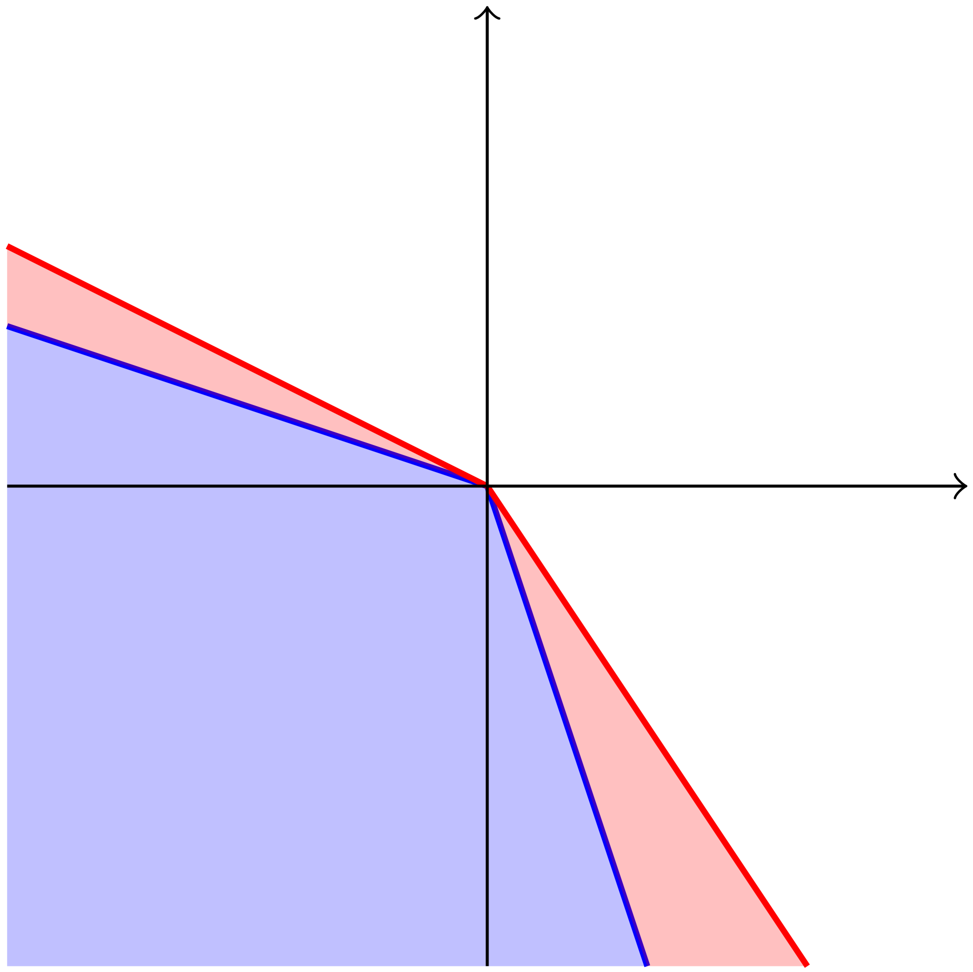

The set filter $\mathscr{C}_{\mathrm{C}}(\omega)$

Fix any ω in Ω.

Non-emptiness

$[0,1] \in \mathscr{C}_{\mathrm{C}}(\omega)$, so any sequence of outcomes ω has at least one stationary forecast that makes it computably random: $\mathscr{C}_{\mathrm{C}}(\omega) \neq \emptyset$.

Increasingness

For all $I, J \in \mathscr{C}$: if $I \in \mathscr{C}_{\mathrm{C}}(\omega)$ and $I \subseteq J$, then $J \in \mathscr{C}_{\mathrm{C}}(\omega)$.

Intersection

For any two interval forecasts I and J in $\mathscr{C}_{\mathrm{C}}(\omega)$, we have that $I \cap J \neq \emptyset$ and $I \cap J \in \mathscr{C}_{\mathrm{C}}(\omega)$.

Randomness is inherently imprecise

Fix any ω in Ω.

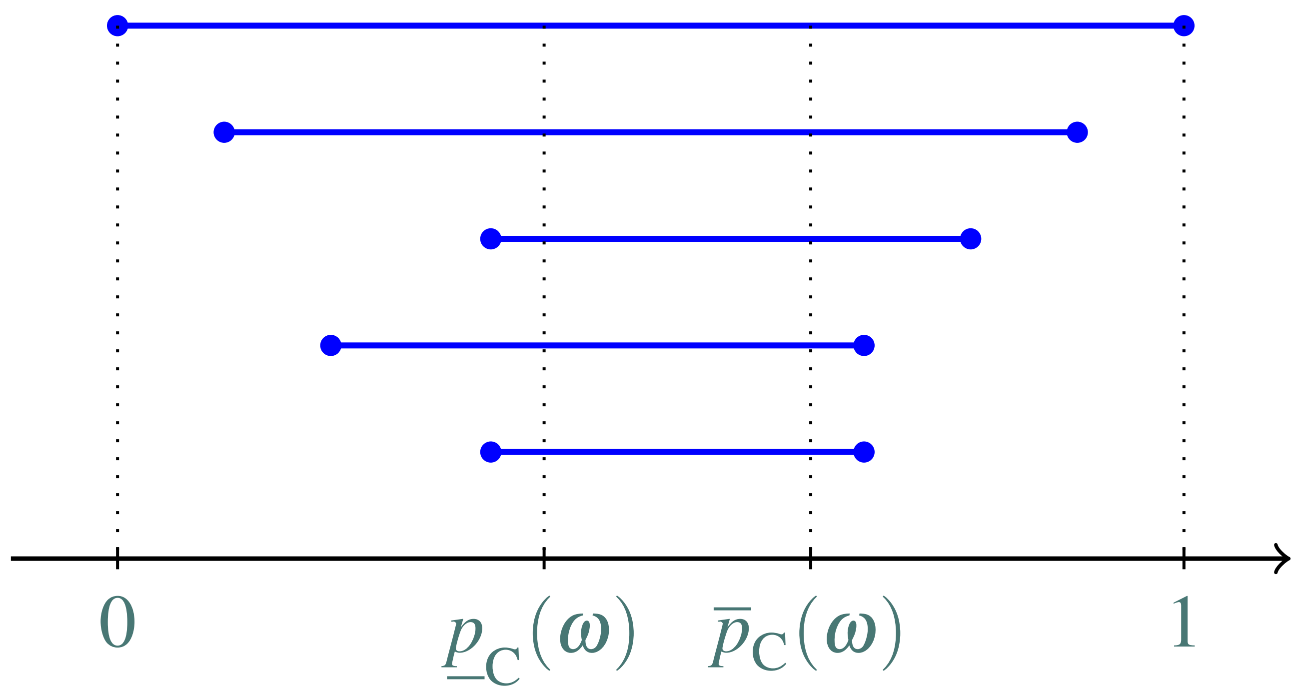

$$\emptyset \neq \bigcap \mathscr{C}_{\mathrm{C}}(\omega) = \left[\underline{p}_{\mathrm{C}}(\omega), \overline{p}_{\mathrm{C}}(\omega)\right].$$

If ω is computable with infinitely many 0's and 1's, then $\left[\underline{p}_{\mathrm{C}}(\omega), \overline{p}_{\mathrm{C}}(\omega)\right] = [0,1]$.

If $\gamma_{\{p\}}$ makes ω computably random, then $\left[\underline{p}_{\mathrm{C}}(\omega), \overline{p}_{\mathrm{C}}(\omega)\right] = \{p\}$.

EXAMPLES

A simple example

Consider any p and q in $[0,1]$ with $p \leq q$, and the forecasting system $\gamma_{p,q}$ defined by

$$\gamma_{p,q}(z_1,\dots,z_n) := \begin{cases} p & \text{if } n \text{ is odd} \\ q & \text{if } n \text{ is even} \end{cases} \quad \text{for all } (z_1,\dots,z_n) \in \Omega^{\Diamond}.$$

Theorem

Consider any outcome sequence ω that is computably random for $\gamma_{p,q}$. Then for all $I \in \mathscr{C}$:

$$I \in \mathscr{C}_{\mathrm{C}}(\omega) \Leftrightarrow [p,q] \subseteq I,$$

and therefore

$$\underline{p}_{\mathrm{C}}(\omega) = p \text{ and } \overline{p}_{\mathrm{C}}(\omega) = q.$$

A more complicated example

$$p_n := \frac{1}{2} + (-1)^n \delta_n, \text{ with } \delta_n := e^{-\frac{1}{n+1}} \sqrt{e^{\frac{1}{n+1}} - 1} \text{ for all } n \in \mathbb{N},$$

Consider the precise forecasting system $\gamma_{\sim 1/2}$ defined by

$$\gamma_{\sim 1/2}(z_1, \dots, z_{n-1}) := p_n \text{ for all } n \in \mathbb{N} \text{ and } (z_1, \dots, z_{n-1}) \in \Omega^{\Diamond}.$$

Theorem

Consider any outcome sequence ω that is computably random for $\gamma_{\sim 1/2}$. Then for all $I \in \mathscr{C}$:

$$I \in \mathscr{C}_{\mathrm{C}}(\omega) \Leftrightarrow \min I < \frac{1}{2} \text{ and } \max I > \frac{1}{2}$$

and therefore

$$\underline{p}_{\mathrm{C}}(\omega) = \overline{p}_{\mathrm{C}}(\omega) = \frac{1}{2}.$$

CONCLUSIONS?


© Tom Roeleveld (CC BY-NC 2.0)
There’s more to uncertainty than probabilities
ISIPTA 2019
3 - 6 July
Ghent, Belgium
The 20-year anniversary edition of the
world’s main forum on imprecise probabilities