Game-theoretic probability

Glenn Shafer, Rutgers University

Mathematics: Game theory in place of measure theory

Interpretation: Validity for probabilities or other prices means only that a speculator will not multiply the capital he risks by a large factor.

Applications:

- Parsimonious explanation of $\sqrt{dt}$, CAPM, & lead-lag effects.
- Defensive forecasting.
- Betting interpretation of Dempster-Shafer.

References

- *Probability and Finance: It's Only a Game!* Glenn Shafer and Vladimir Vovk, Wiley, 2001.

- www.probabilityandfinance.com: Chapters from book, reviews, many working papers.

Heroes of game-theoretic probability

Blaise Pascal
(1623–1662)

Probability is about betting.

Antoine Cournot
(1801–1877)

Events of small probability do not happen.

Jean Ville
(1910–1988)

Pascal + Cournot:

If the probabilities are right, you don't get rich.

Antoine Cournot
(1801–1877)

"A physically impossible event is one whose probability is infinitely small. This remark alone gives substance—an objective and phenomenological value—to the mathematical theory of probability." (1843)

This is more basic than frequentism.

Émile Borel
1871–1956

Inventor of measure theory.

Minister of the French navy in 1925.

Borel was emphatic: the principle that an event with very small probability will not happen is the only law of chance.

- Impossibility on the human scale: $p < 10^{-6}$.
- Impossibility on the terrestrial scale: $p < 10^{-15}$.
- Impossibility on the cosmic scale: $p < 10^{-50}$.

Andrei Kolmogorov
1903–1987

Hailed as the Soviet Euler, Kolmogorov was credited with establishing measure theory as the mathematical foundation for probability.

In his celebrated 1933 book, Kolmogorov wrote:

> When P(A) very small, we can be practically certain that the event A will not happen on a single trial of the conditions that define it.

Jean Ville,
1910–1988, on
entering the *École
Normale Supérieure.*

In 1939, Ville showed that the laws of probability can be derived from this principle:

> You will not multiply the capital you risk by a large factor.

Ville showed that this principle is equivalent to the principle that events of small probability will not happen.

We call both principles Cournot's principle.

Suppose you gamble without risking more than your initial capital.

Your resulting wealth is a nonnegative random variable X with expected value $E(X)$ equal to your initial capital.

Markov's inequality says

$$P\left(X \geq \frac{E(X)}{\epsilon}\right) \leq \epsilon.$$

You have probability ϵ or less of multiplying your initial capital by $1/\epsilon$ or more.

Ville proved what is now called *Doob's inequality*, which generalizes Markov's inequality to a sequence of bets.

Volodya Vovk atop the World Trade Center in 1998.

- Born 1960.
- Student of Kolmogorov.
- Born in Ukraine, educated in Moscow, teaches in London.
- Volodya is a nickname for the Ukrainian Volodimir and the Russian Vladimir.

The Ville/Vovk perfect-information protocol for probability

$\mathcal{K}_0 = 1$.
FOR $n = 1, 2, \ldots, N$:
    Forecaster announces prices for various payoffs.
    Skeptic decides which payoffs to buy.
    Reality determines the payoffs.
    $\mathcal{K}_n := \mathcal{K}_{n-1}$ + Skeptic's net gain or loss.

Ville showed that any test of Forecaster can be expressed as a betting strategy for Skeptic.

Vovk, Takemura, and I showed that Forecaster can beat Skeptic.

Ville/Vovk game-theoretic testing

In Ville's theory, Forecaster is a known probability distribution for Reality's move. It always gives conditional probabilities for Reality's next move given her past moves.

In Vovk's generalization, (1) Forecaster does not necessarily use a known probability distribution, and (2) he may give less than a probability distribution for Reality's next move. For both reasons, we get upper and lower probabilities instead of probabilities.

Ville's strong law of large numbers.

(Special case where probability is always 1/2.)

$\mathcal{K}_0 = 1$.
FOR $n = 1, 2, \ldots$:
Skeptic announces $s_n \in \mathbb{R}$.
Reality announces $y_n \in \{0, 1\}$.
$\mathcal{K}_n := \mathcal{K}_{n-1} + s_n(y_n - \frac{1}{2})$.

Skeptic wins if
(1) $\mathcal{K}_n$ is never negative and
(2) either $\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} y_i = \frac{1}{2}$ or $\lim_{n\to\infty} \mathcal{K}_n = \infty$.

Theorem Skeptic has a winning strategy.

Ville's strategy

> $\mathcal{K}_0 = 1$.
> FOR $n = 1, 2, \ldots$:
>   Skeptic announces $s_n \in \mathbb{R}$.
>   Reality announces $y_n \in \{0, 1\}$.
> $\mathcal{K}_n := \mathcal{K}_{n-1} + s_n(y_n - \frac{1}{2})$.

Ville suggested the strategy

$$s_n(y_1, \ldots, y_{n-1} = \frac{4}{n+1}\mathcal{K}_{n-1}\left(r_{n-1} - \frac{n-1}{2}\right), \text{ where } r_{n-1} := \sum_{i=1}^{n-1} y_i.$$

It produces the capital

$$\mathcal{K}_n = 2^n \frac{r_n!(n-r_n)!}{(n+1)!}.$$

From the assumption that this remains bounded by some constant C, you can easily derive the strong law of large numbers using Stirling's formula.

Vovk's weak law of large numbers

$\mathcal{K}_0 := 1$.
FOR $n = 1, \dots, N$:
  Skeptic announces $s_n \in \mathbb{R}$.
  Reality announces $y_n \in \{0, 1\}$.
  $\mathcal{K}_n := \mathcal{K}_{n-1} + s_n \left(y_n - \frac{1}{2}\right)$.

Winning: Skeptic wins if $\mathcal{K}_n$ is never negative and either $\mathcal{K}_N \geq C$ or $|\frac{1}{N}\sum_{n=1}^{N} y_n - \frac{1}{2}| < \epsilon$.

Theorem. Skeptic has a winning strategy if $N \geq C/4\epsilon^2$.

Ville's more general game.

Ville started with a probability distribution for P for $y_1, y_2, \ldots$.
The conditional probability for $y_n = 1$ given $y_1, \ldots, y_{n-1}$ is not necessarily $1/2$.

$\mathcal{K}_0 := 1$.
FOR $n = 1, 2, \ldots$:
Skeptic announces $s_n \in \mathbb{R}$.
Reality announces $y_n \in \{0, 1\}$.
$\mathcal{K}_n := \mathcal{K}_{n-1} + s_n \left(y_n - \mathrm{P}(y_n = 1 | y_1, \ldots, y_{n-1})\right)$.
Skeptic wins if
(1) $\mathcal{K}_n$ is never negative and
(2) either $\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} \left(y_i - \mathrm{P}(y_i = 1 | y_1, \ldots, y_{i-1})\right) = 0$
or $\lim_{n\to\infty} \mathcal{K}_n = \infty$.

Theorem Skeptic has a winning strategy.

Vovk's generalization: Replace P with a forecaster.

$\mathcal{K}_0 := 1$.

FOR $n = 1, 2, \ldots$:

Forecaster announces $p_n \in [0, 1]$.

Skeptic announces $s_n \in \mathbb{R}$.

Reality announces $y_n \in \{0, 1\}$.

$\mathcal{K}_n := \mathcal{K}_{n-1} + s_n(y_n - p_n)$.

Skeptic wins if

(1) $\mathcal{K}_n$ is never negative and

(2) either $\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} (y_i - p_i) = 0$

or $\lim_{n\to\infty} \mathcal{K}_n = \infty$.

Theorem Skeptic has a winning strategy.

Vovk's weak law of large numbers

$\mathcal{K}_0 := 1$.
FOR $n = 1, \ldots, N$:

Forecaster announces $p_n \in [0,1]$.
Skeptic announces $s_n \in \mathbb{R}$.
Reality announces $y_n \in \{0,1\}$.
$\mathcal{K}_n := \mathcal{K}_{n-1} + s_n (y_n - p_n)$.

Winning: Skeptic wins if $\mathcal{K}_n$ is never negative and either $\mathcal{K}_N \geq C$ or $\left|\frac{1}{N}\sum_{n=1}^{N}(y_n - p_n)\right| < \epsilon$.

Theorem. Skeptic has a winning strategy if $N \geq C/4\epsilon^2$.

Put it in terms of upper probability

$\mathcal{K}_0 := 1$.
FOR $n = 1, \dots, N$:
  Forecaster announces $p_n \in [0,1]$.
  Skeptic announces $s_n \in \mathbb{R}$.
  Reality announces $y_n \in \{0,1\}$.
  $\mathcal{K}_n := \mathcal{K}_{n-1} + s_n(y_n - p_n)$.

Theorem. $\overline{\mathbb{P}}\left\{\frac{1}{N}|\sum_{n=1}^{N}(y_n - p_n)| \geq \epsilon\right\} \leq \frac{1}{4N\epsilon^2}$.

Definition of upper price and upper probability

$\mathcal{K}_0 := \alpha$.
FOR $n = 1, \ldots, N$:
  Forecaster announces $p_n \in [0,1]$.
  Skeptic announces $s_n \in \mathbb{R}$.
  Reality announces $y_n \in \{0,1\}$.
  $\mathcal{K}_n := \mathcal{K}_{n-1} + s_n (y_n - p_n)$.

For any real-valued function X on $([0,1] \times \{0,1\})^N$,

$\overline{\mathbb{E}}\, X := \inf\{\alpha \,|\, \text{Skeptic has a strategy guaranteeing } \mathcal{K}_N \geq X(p_1, y_1, \ldots, p_N, y_N)\}$

For any subset $A \subseteq ([0,1] \times \{0,1\})^N$,

$$\overline{\mathbb{P}}\, A := \inf\{\alpha \,|\, \text{Skeptic has a strategy guaranteeing } \mathcal{K}_N \geq 1 \text{ if } A \text{ happens and } \mathcal{K}_N \geq 0 \text{ otherwise}\}.$$

$$\underline{\mathbb{E}}\, X = -\overline{\mathbb{E}}(-X) \qquad\qquad \underline{\mathbb{P}}\, A = 1 - \overline{\mathbb{P}}\, \overline{A}$$

Defensive forecasting

Under repetition, good probability forecasting is possible.

- We call it defensive because it defends against a quasi-universal test.

- Your probability forecasts will pass this test even if reality plays against you.

Why Phil Dawid thought good probability prediction is impossible. . .

FOR $n = 1, 2, \ldots$
  Forecaster announces $p_n \in [0, 1]$.
  Skeptic announces continuous $s_n \in \mathbb{R}$.
  Reality announces $y_n \in \{0, 1\}$.
  Skeptic's profit $:= s_n(y_n - p_n)$.

Reality can make Forecaster uncalibrated by setting

$$y_n := \begin{cases} 1 & \text{if } p_n < 0.5 \\ 0 & \text{if } p_n \geq 0.5, \end{cases}$$

Skeptic can then make steady money with

$$s_n := \begin{cases} 1 & \text{if } p < 0.5 \\ -1 & \text{if } p \geq 0.5, \end{cases}$$

But if Skeptic is forced to approximate s_n by a continuous function of p_n, then the continuous function will be zero close to $p = 0.5$, and Forecaster can set p_n equal to this point.

Skeptic adopts a continuous strategy $\mathcal{S}$.
FOR $n = 1, 2, \ldots$
Reality announces $x_n \in \mathbf{X}$.
Forecaster announces $p_n \in [0, 1]$.
Skeptic makes the move s_n specified by $\mathcal{S}$.
Reality announces $y_n \in \{0, 1\}$.
Skeptic's profit $:= s_n(y_n - p_n)$.

Theorem Forecaster can guarantee that Skeptic never makes money.

We actually prove a stronger theorem. Instead of making Skeptic announce his entire strategy in advance, only make him reveal his strategy for each round in advance of Forecaster's move.

FOR $n = 1, 2, \ldots$
Reality announces $x_n \in \mathbf{X}$.
Skeptic announces continuous $S_n : [0, 1] \to \mathbb{R}$.
Forecaster announces $p_n \in [0, 1]$.
Reality announces $y_n \in \{0, 1\}$.
Skeptic's profit $:= S_n(p_n)(y_n - p_n)$.

Theorem. Forecaster can guarantee that Skeptic never makes money.

FOR $n = 1, 2, \ldots$
Reality announces $x_n \in \mathbf{X}$.
Skeptic announces continuous $S_n : [0, 1] \to \mathbb{R}$.
Forecaster announces $p_n \in [0, 1]$.
Reality announces $y_n \in \{0, 1\}$.
Skeptic's profit $:= S_n(p_n)(y_n - p_n)$.

Theorem Forecaster can guarantee that Skeptic never makes money.

Proof:

- If $S_n(p) > 0$ for all p, take $p_n := 1$.
- If $S_n(p) < 0$ for all p, take $p_n := 0$.
- Otherwise, choose p_n so that $S_n(p_n) = 0$.

TWO APPROACHES TO FORECASTING

FOR $n = 1, 2, \ldots$
  Forecaster announces $p_n \in [0, 1]$.
  Skeptic announces $s_n \in \mathbb{R}$.
  Reality announces $y_n \in \{0, 1\}$.

1. Start with strategies for Forecaster. Improve by averaging (Bayes, prediction with expert advice).

2. Start with strategies for Skeptic. Improve by averaging (defensive forecasting).

We can always give probabilities with good calibration and resolution.

FOR $n = 1, 2, \dots$ Forecaster announces $p_n \in [0, 1]$. Reality announces $y_n \in \{0, 1\}$.

There exists a strategy for Forecaster that gives p_n with good calibration and resolution.

FOR $n = 1, 2, \dots$
  Reality announces $x_n \in \mathbf{X}$.
  Forecaster announces $p_n \in [0, 1]$.
  Reality announces $y_n \in \{0, 1\}$.

1. Fix $p^* \in [0, 1]$. Look at n for which $p_n \approx p^*$. If the frequency of $y_n = 1$ always approximates p^*, Forecaster is *properly calibrated*.

2. Fix $x^* \in \mathbf{X}$ and $p^* \in [0, 1]$. Look at n for which $x_n \approx x^*$ and $p_n \approx p^*$. If the frequency of $y_n = 1$ always approximates p^*, Forecaster is properly calibrated and has *good resolution*.

Fundamental idea: Average strategies for Skeptic for a grid of values of p^*. (The p^*-strategy makes money if calibration fails for p_n close to p^*.) The derived strategy for Forecaster guarantees good calibration everywhere.

Example of a resulting strategy for Skeptic:

$$S_n(p) := \sum_{i=1}^{n-1} e^{-C(p-p_i)^2}(y_i - p_i)$$

Any kernel $K(p, p_i)$ can be used in place of $e^{-C(p-p_i)^2}$.

Skeptic's strategy:

$$S_n(p) := \sum_{i=1}^{n-1} e^{-C(p-p_i)^2}(y_i - p_i)$$

Forecaster's strategy: Choose p_n so that

$$\sum_{i=1}^{n-1} e^{-C(p_n-p_i)^2}(y_i - p_i) = 0.$$

The main contribution to the sum comes from i for which p_i is close to p_n. So Forecaster chooses p_n in the region where the $y_i - p_i$ average close to zero.

On each round, choose as p_n the probability value where calibration is the best so far.

The $\sqrt{dt}$ effect:

- The average change over one day is about 22% of the average change over one month. ($\sqrt{1/20} \approx 0.22$)

- The average change over one day is about 6% of the average change over one year. ($\sqrt{1/250} \approx 0.06$)

- The average change over one year is about 32% of the average change over ten years. ($\sqrt{1/10} \approx 0.32$)

Why does the $\sqrt{dt}$ effect happen?

Because otherwise a speculator could multiply the capital he risks by a large factor.

- If prices are more jagged than $\sqrt{dt}$ (daily changes tend to exceed 6% of annual changes), then a simple contrarian strategy can make a lot of money.

- If prices are less jagged than $\sqrt{dt}$ (daily changes tend to be less than 6% of annual changes), then a simple momentum strategy can make a lot of money.

- More jagged than $\sqrt{dt}$ means $\sum_n |dS_n|^2$ is large relative to $\max_n |S_n - S_0|$.

- Less jagged than $\sqrt{dt}$ means $\sum_n |dS_n|^2$ is small relative to $\max_n |S_n - S_0|$.

- Less jagged than $\sqrt{dt}$ means $\sum_n |dS_n|^2$ is small relative to $\max_n |S_n - S_0|$.

- If we can count on $\sum_n |dS_n|^2 \leq \sigma^2_{\max}$ and $\max_n |S_n - S_0| \geq D$, then a simple momentum strategy can turn \$1 into $\$D^2/\sigma^2_{\max}$ or more for sure.

- To wit, invest

$$2\frac{1}{\sigma^2_{\max}}S_{n-1}$$

in the security on round n.

References

- La variation d'ordre p des semi-martingales, by Dominique Lepingle, *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 36, 295–316, 1976.

- Continuous price processes in frictionless markets have infinite variation, by J. Michael Harrison, Richard Pitbladdo, and Stephan M. Schaefer, *Journal of Business*, 57:3, 353–365, 1984.

- Arbitrage with fractional Brownian motion, by L. Chris G. Rogers, *Mathematical Finance*, 7, 95–105, 1997.

- Arbitrage in fractional Brownian motion models, by Patrick Cheridito, *Finance and Stochastics*, 7:4, 533-553, October, 2003.

CLASSICAL CAPM

$$\mathrm{E}(\tilde{s}) = f + (\mathrm{E}(\tilde{m}) - f)\frac{\mathrm{Cov}(\tilde{s}, \tilde{m})}{\mathrm{Var}(\tilde{m})}.$$

- $\tilde{s}$ is the random variable whose realization is the simple return s for the stock.
- $\tilde{m}$ is the random variable whose realization is the simple return m for the market index.
- f is risk-free rate.

The game-theoretic CAPM is an analogous relation between empirical (ex post) quantities:

$$\mu_s := \frac{1}{N}\sum_{n=1}^{N} s_n \qquad \mu_m := \frac{1}{N}\sum_{n=1}^{N} m_n$$

$$\sigma_m^2 := \frac{1}{N}\sum_{n=1}^{N} m_n^2 \qquad \sigma_{sm} := \frac{1}{N}\sum_{n=1}^{N} s_n m_n$$

$$\beta_s := \sigma_{sm}/\sigma_m^2$$

GAME-THEORETIC CAPM

$$\mu_s \approx (\mu_m - \sigma_m^2) + \sigma_m^2 \beta_s, \tag{1}$$

If we write μ_f for $\mu_m - \sigma_m^2$, then the game-theoretic CAPM can be written in the form

$$\mu_s \approx \mu_f + (\mu_m - \mu_f)\beta_s.$$

Aleatory (objective) vs. epistemic (subjective)

From a 1970s perspective:

- Aleatory probability is the irreducible uncertainty that remains when knowledge is complete.

- Epistemic probability arises when knowledge is incomplete.

New game-theoretic perspective:

- Under a repetitive structure you can make make good probability forecasts relative to whatever state of knowledge you have.

- If there is no repetitive structure, your task is to combine evidence rather than to make probability forecasts.

Cournotian understanding of Dempster-Shafer

- Fundamental idea: transferring belief
- Conditioning
- Independence
- Dempster's rule

Fundamental idea: transferring belief

- Variable ω with set of possible values Ω.
- Random variable $\mathbf{X}$ with set of possible values $\mathcal{X}$.
- We learn a mapping $\Gamma : \mathcal{X} \to 2^{\Omega}$ with this meaning:

$$\text{If } \mathbf{X} = x, \text{ then } \omega \in \Gamma(x).$$

- For $A \subseteq \Omega$, our belief that $\omega \in A$ is now

$$\mathbb{B}(A) = \mathbb{P}\{x | \Gamma(x) \subseteq A\}.$$

Cournotian judgement of independence: Learning the relationship between $\mathbf{X}$ and ω does not affect our inability to beat the probabilities for $\mathbf{X}$.

Example: The sometimes reliable witness

- Joe is reliable with probability 30%. When he is reliable, what he says is true. Otherwise, it may or may not be true.

  $\mathcal{X} = \{\text{reliable}, \text{not reliable}\} \qquad \mathbb{P}(\text{reliable}) = 0.3 \qquad \mathbb{P}(\text{not reliable}) = 0.7$

- Did Glenn pay his dues for coffee? $\qquad \Omega = \{\text{paid}, \text{not paid}\}$

- Joe says "Glenn paid."

  $$\Gamma(\text{reliable}) = \{\text{paid}\} \qquad \Gamma(\text{not reliable}) = \{\text{paid}, \text{not paid}\}$$

- New beliefs:

  $$\mathbb{B}(\text{paid}) = 0.3 \qquad \mathbb{B}(\text{not paid}) = 0$$

Cournotian judgement of independence: Hearing what Joe said does not affect our inability to beat the probabilities concerning his reliability.

Art Dempster (born 1929) with his Meng & Shafer hatbox.

Retirement dinner at Harvard, May 2005.

See http://www.stat.purdue.edu/ chuanhai/projects/DS/ for Art's D-S papers.