Is everything stochastic?

Glenn Shafer
Rutgers University

Games and Decisions
Centro di Ricerca Matematica Ennio De Giorgi
8 July 2013

1. Game theoretic probability
2. Game theoretic upper and lower probability
3. Defensive forecasting
4. Is everything stochastic?

Preview

1. Game theoretic probability
2. Game theoretic upper and lower probability
3. Defensive forecasting
4. Is everything stochastic?

1. Preview: Game theoretic probability
2. Preview: Game theoretic upper and lower probabilities
3. Preview: Defensive forecasting
4. Preview: Is everything stochastic?

- **Probabilities derive from betting offers.**

  Not from the measure of sets.

- **Probabilities are tested by betting strategies.**

- **Probability theorems are proven by betting strategies.**

  -- Do not say that the property fails on a set of measure zero.

  -- Say that its failure implies the success of a betting strategy.

1. Preview: Game theoretic probability
2. Preview: Game theoretic upper and lower probabilities
3. Preview: Defensive forecasting
4. Preview: Is everything stochastic?

Probabilities derive from betting offers.

The offers may determine less than a probability distribution.

1. The stock market does not give a probability distribution for tomorrow's price of share of PotashCorp.

2. The weather forecaster who gives probabilities for rain the next day over an entire year does not give a joint probability distribution for the 365 outcomes.

In such cases, we get only upper and lower probabilities.

1. Preview: Game theoretic probability
2. Preview: Game theoretic upper and lower probability
3. Preview: Defensive forecasting
4. Preview: Is everything stochastic?

- In the game-theoretic framework, it can be shown that **good probability forecasting is possible**.

- For a sequence of events, **you can give step-by-step probabilities that pass statistical tests**.

- The forecasting **defends** against the tests.

Step-by-Step

Assumptions:

1. **You give a probability for each event in a sequence.**

2. **You see the prior outcomes before you give each new probability.**

Under these assumptions, you can choose the probabilities so they pass statistical tests.

- Modeling is not needed.

- The sequence need not be “iid”; this concept is not even defined.

1. Preview: Game theoretic probability
2. Preview: Game theoretic upper and lower probability
3. Preview: Defensive forecasting
4. Preview: Is everything stochastic?

Jeyzy Neyman's **inductive behavior**

A statistician who makes predictions with 95% confidence has two goals:

be informative

be right 95% of the time

Question: Why isn't this good enough for **probability judgment**?

Answer: Because two statisticians who are right 95% of the time may tell the court different and even contradictory things.

They are placing the current event in different sequences.

- Good probability forecasting requires a sequence,
- It does not require necessarily repetition of the "same" event.
- Each event remains unique.

Probability judgment: Assessment of the relevance or irrelevance of information to the ability of a probability forecaster to defeat tests in a given sequence.

The game-theoretic framework for probability

1. Pascal’s game-theoretic probability
2. Making Pascal’s game precise
3. Game-theoretic testing and Cournot’s principle
4. Defensive forecasting
5. Is everything stochastic?

1. Pascal’s game-theoretic probability

The contrast between measure-theoretic & game-theoretic probability began in 1654.

Pascal = game theory

Fermat = measure theory

Pascal's question to Fermat in 1654

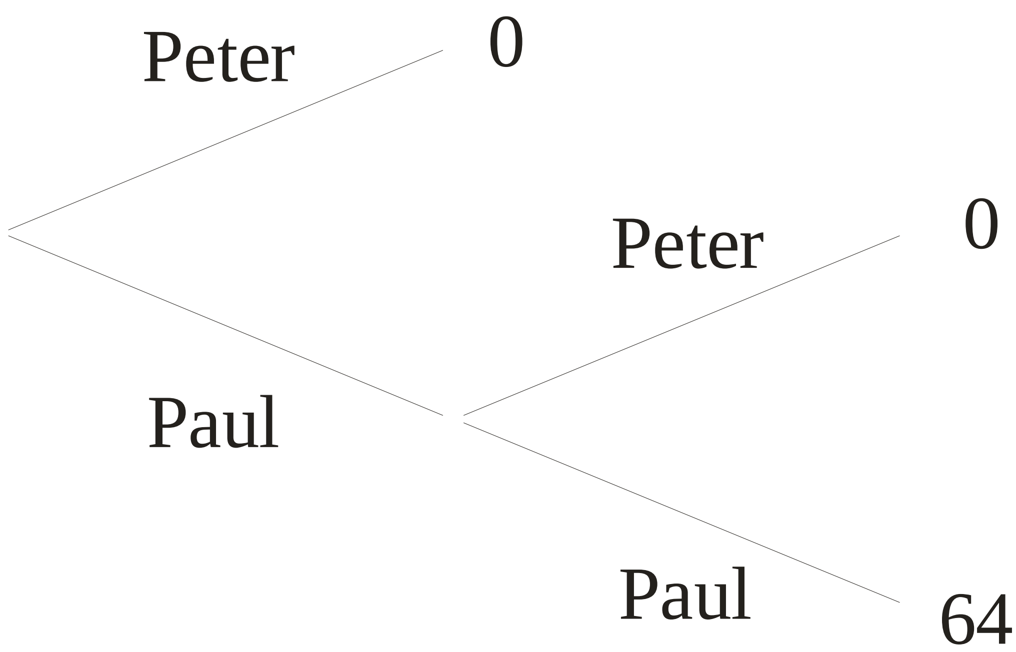

Blaise Pascal (1623–1662).

Paul needs 2 points to win.
Peter needs only one.

If the game must be broken off, how much of the stake should Paul get?

Fermat's answer (measure theory)

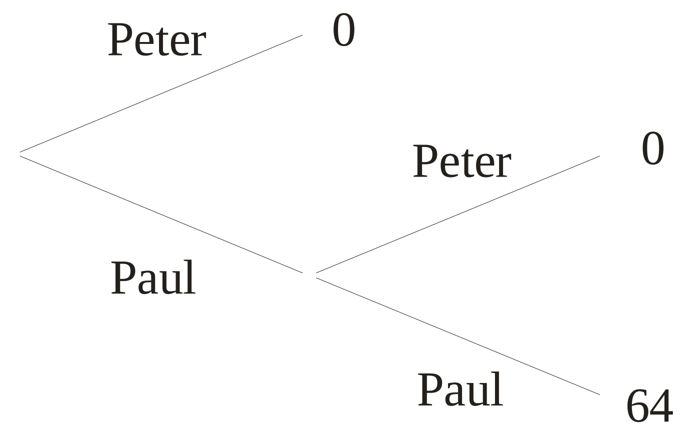

Count the possible outcomes.

Suppose they play two rounds.   There are 4 possible outcomes:

1. Peter wins first, Peter wins second
2. Peter wins first, Paul wins second
3. Paul wins first, Peter wins second
4. Paul wins first, Paul wins second

Paul wins only in outcome 4. So his share should be ¼, or 16 pistoles.

Pascal didn't like the argument.

Pierre Fermat, 1601-1665

Pascal's answer (game theory)

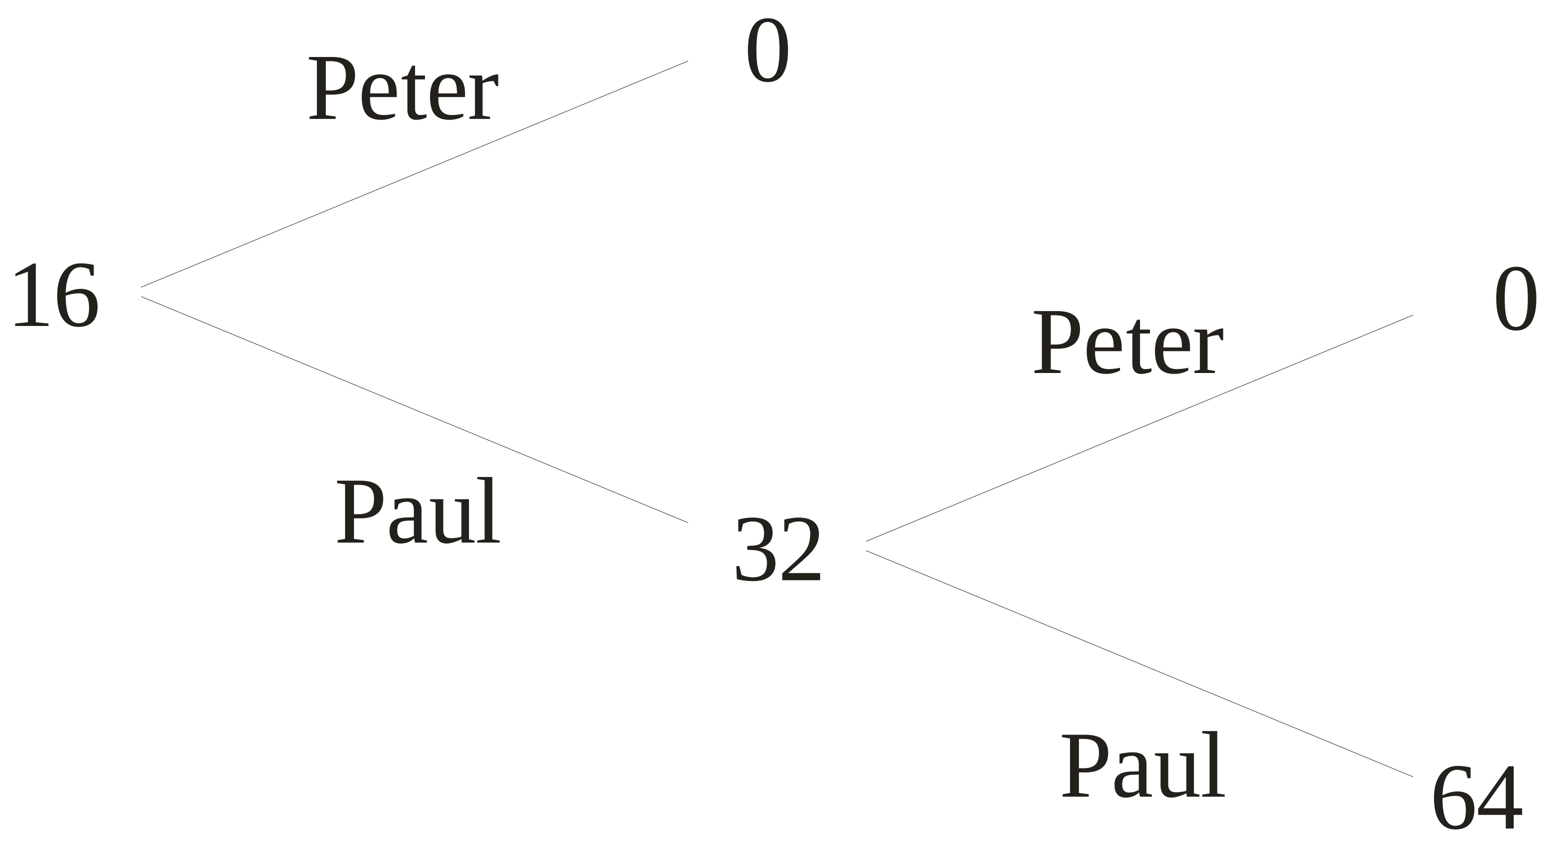

Measure-theoretic probability:
- Classical: elementary events with probabilities adding to one.
- Modern: space with filtration and probability measure.

Probability of A = total of probabilities for elementary events favoring A

Game-theoretic probability:
- One player offers prices for uncertain payoffs.
- Another player decides what to buy.

Probability of A = initial stake needed to obtain the payoff
[1 if A happens and 0 otherwise]

If no strategy delivers exactly the 0/1 payoff:

Upper probability of A = initial stake needed to obtain <u>at least</u> the payoff
[1 if A happens, 0 otherwise]

The game-theoretic framework for probability

1. Pascal's game-theoretic probability
2. Making Pascal's game precise
3. Game-theoretic testing and Cournot's principle
4. Defensive forecasting
5. Is everything stochastic?

To make Pascal’s theory part of modern game theory, we must define the game precisely.

- Rules of play
- Each player’s information
- Rule for winning

A game between Forecaster and Reality

Forecaster gives probabilities for a sequence $x_1, x_2, \ldots$ of 1s and 0s.

Before Reality announces x_n, Forecaster announces probability p_n for $x_n = 1$.

FOR $n = 1, 2, \ldots$:
  Forecaster announces $p_n \in [0, 1]$.
  Reality announces $x_n \in \{0, 1\}$.

Theorem: Forecaster can give $p_1, p_2, \ldots$ that are not refuted by statistical tests.

FOR $n = 1, 2, \ldots$:
Forecaster announces $p_n \in [0, 1]$.
Reality announces $x_n \in \{0, 1\}$.

Clarifications:

1. The phenomena need not be binary. We assume $x_n \in \{0, 1\}$ only for simplicity.
2. Reality's move space may change from round to round.
3. Perfect information: All players hear announcements as they are made.
4. In addition to $x_1, \ldots, x_{n-1}$, Forecaster may have other newly acquired information.
5. To be fair to Forecaster, we do not consider statistical tests based on information he does not have.

Forecaster with a strategy P

Suppose Reality plays

$$x_1 = 1$$
$$x_2 = 0$$
$$\ldots$$

If Forecaster begins the game with a probability distribution P for $x_1, x_2, \ldots$, then he can set

$$p_1 := \mathrm{P}(x_1 = 1)$$
$$p_2 := \mathrm{P}(x_2 = 1 | x_1 = 1)$$
$$p_3 := \mathrm{P}(x_3 = 1 | x_1 = 1 \;\; \& \;\; x_2 = 0)$$
$$\ldots$$

Alternatively, if Forecaster begins the game with a probability distribution P for everything he might see as the game proceeds, then he can set

$$p_1 := \mathrm{P}(x_1 = 1 | \text{all info before round 1})$$
$$p_2 := \mathrm{P}(x_2 = 1 | \text{all info before round 2})$$
$$p_3 := \mathrm{P}(x_3 = 1 | \text{all info before round 3})$$
$$\ldots$$

But Forecaster is not required to base his moves on an initial probability distribution P.

FOR $n = 1, 2, \ldots$:
  Forecaster announces $p_n \in [0,1]$.
  Reality announces $x_n \in \{0,1\}$.

Forecaster's moves do not define a probability distribution for $x_1, x_2, \ldots$.

If Reality plays $x_1 = 1$, $x_2 = 0$, and so on, then Forecaster's moves $p_1, p_2, \ldots$ can be interpreted as conditional probabilities:

$$\begin{aligned}
p_1 &= \mathrm{P}(x_1 = 1) \\
p_2 &= \mathrm{P}(x_2 = 1 | x_1 = 1) \\
p_3 &= \mathrm{P}(x_3 = 1 | x_1 = 1 \ \ \& \ \ x_2 = 0) \\
&\ldots
\end{aligned}$$

But these conditional probabilities fall short of defining a probability distribution P for $x_1, x_2, \ldots$. They leave unspecified the conditional probabilities

$$\begin{aligned}
&\mathrm{P}(x_2 = 1 | x_1 = 0) \\
&\mathrm{P}(x_3 = 1 | x_1 = 0 \ \ \& \ \ x_2 = 0) \\
&\mathrm{P}(x_3 = 1 | x_1 = 0 \ \ \& \ \ x_2 = 1) \\
&\mathrm{P}(x_3 = 1 | x_1 = 1 \ \ \& \ \ x_2 = 1) \\
&\ldots
\end{aligned}$$

A probability distribution P is a **strategy** for Forecaster, and Forecaster can play without a strategy.

The game-theoretic framework for probability

1. Pascal's game-theoretic probability
2. Making Pascal's game precise
3. Game-theoretic testing and Cournot's principle
4. Defensive forecasting
5. Is everything stochastic?

Forecaster is tested by a third player, Skeptic, who tries to get rich from Forecaster's betting offers.

Players: Forecaster, Reality, Skeptic

Protocol:

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

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

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

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

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

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

Winner: Skeptic wins if $\mathcal{K}_n \geq 0$ for all n and $\lim_{n\to\infty} \mathcal{K}_n = \infty$. Otherwise Forecaster and Reality win.

Example of a game-theoretic probability theorem.

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

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.

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.

Winner: Skeptic wins if $\mathcal{K}_n \geq 0$ for all n and $\lim_{n\to\infty} \mathcal{K}_n = \infty$. Otherwise Forecaster and Reality win.

The thesis that statistical testing can be always be carried out by <u>strategies that attempt to multiply the capital risked</u> goes back to Ville.

Jean André Ville, 1910-1989

At home at 3, rue Campagne Première, shortly after the Liberation

Empirical interpretation of probability

Cournot's principle

Commonly accepted by mathematicians before WWII

An event of very small probability will not happen.

To avoid lottery paradox, consider only events with simplest descriptions.
(Wald, Schnorr, Kolmogorov, Levin)

Ville's principle

Equivalent to Cournot's principle when upper probabilities are probabilities

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

Mathematical definition of probability:

P(A) = stake needed to obtain $1 if A happens, $0 otherwise

Objective (empirical) interpretation of game-theoretic probability:

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

Subjective interpretation of game-theoretic probability:

I don’t think you will multiply the capital you risk by a large factor.

Unlike de Finetti, we do not need behavioral assumptions (e.g., people want to bet or can be forced to do so).

For more on statistical testing by martingales, see my 2001 book with Kolmogorov's student Volodya Vovk.

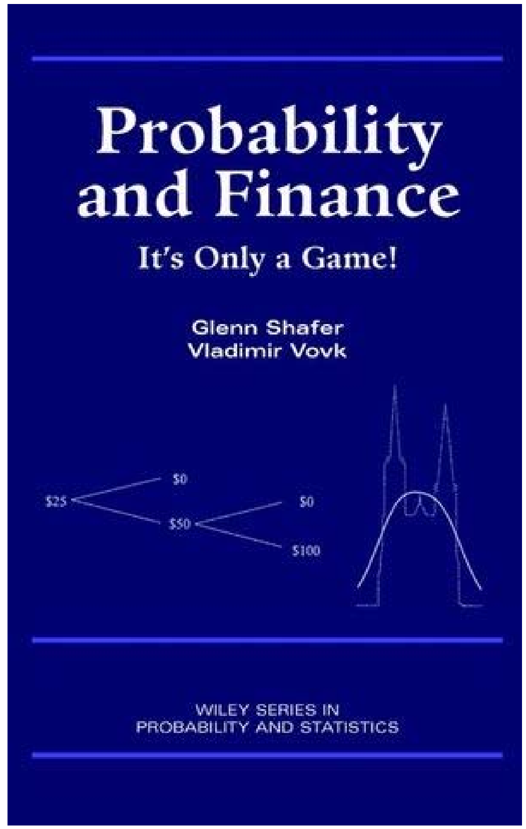

10 years of subsequent working papers at www.probabilityandfinance.com

The game-theoretic framework for probability

1. Pascal's game-theoretic probability
2. Making Pascal's game precise
3. Game-theoretic testing and Cournot's principle
4. Defensive forecasting
5. Is everything stochastic?

Two paths to successful probability forecasting

1. Insist that tests be continuous. Conventional tests can be implemented with continuous betting strategies (Shafer & Vovk, 2001). Only continuous functions are constructive (L. E. J. Brouwer).

Leonid Levin, born 1948

2. Allow Forecaster to hide his precise prediction from Reality using a bit of randomization.

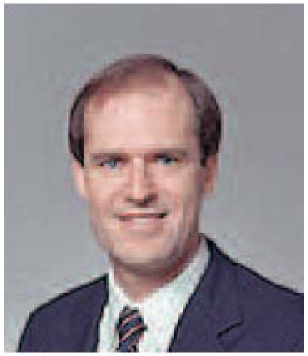

Dean Foster

Rakesh Vohra

Defensive forecasting

The name was introduced in Working Paper 8 at www.probabilityandfinance, by Vovk, Takemura, and Shafer (September 2004). See also Working Papers 7, 9, 10, 11, 13, 14, 16, 17, 18, 20, 21, 22, and 30.

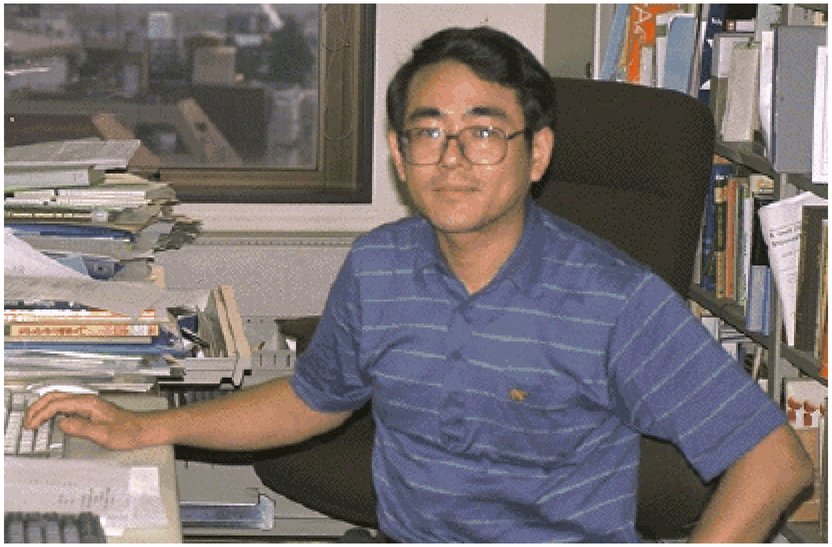

Akimichi Takemura

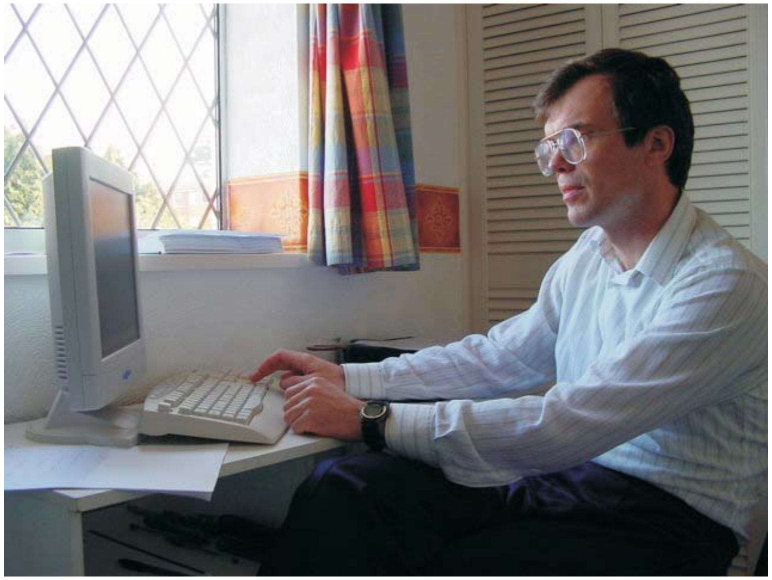

Volodya Vovk

Crucial idea: all the tests (betting strategies for Skeptic) Forecaster needs to pass can be merged into a single portmanteau test for Forecaster to pass.

1. If you have two strategies for multiplying capital risked, divide your capital between them.

2. Formally: average the strategies.

3. You can average countably many strategies.

4. As a practical matter, there are only countably many tests (Abraham Wald, 1937).

5. **I will explain how Forecaster can beat any single test (including the portmanteau test).**

A. How Forecaster beats any single test

B. How to construct a portmanteau test for binary probability forecasting

- Use law of large numbers to test calibration for each probability p.
- Merge the tests for different p.

A game between Forecaster and Reality

Forecaster gives probabilities for a sequence $x_1, x_2, \ldots$ of 1s and 0s.

Before Reality announces x_n, Forecaster announces probability p_n for $x_n = 1$.

FOR $n = 1, 2, \ldots$:
  Forecaster announces $p_n \in [0, 1]$.
  Reality announces $x_n \in \{0, 1\}$.

Theorem: Forecaster can give $p_1, p_2, \ldots$ that are not refuted by statistical tests.

FOR $n = 1, 2, \ldots$:
  Forecaster announces $p_n \in [0, 1]$.
  Reality announces $x_n \in \{0, 1\}$.

Clarifications:

1. The phenomena need not be binary. We assume $x_n \in \{0, 1\}$ only for simplicity.

2. Reality's move space may change from round to round.

3. Perfect information: All players hear announcements as they are made.

4. In addition to $x_1, \ldots, x_{n-1}$, Forecaster may have other newly acquired information.

5. To be fair to Forecaster, we do not consider statistical tests based on information he does not have.

Forecaster is tested by a third player, Skeptic, who tries to get rich from Forecaster's betting offers.

Players: Forecaster, Reality, Skeptic

Protocol:

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

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

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

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

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

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

Winner: Skeptic wins if $\mathcal{K}_n \geq 0$ for all n and $\lim_{n\to\infty} \mathcal{K}_n = \infty$. Otherwise Forecaster and Reality win.

How Forecaster can beat any single test *S*

Skeptic adopts a continuous strategy $\mathcal{S}$.
FOR $n = 1, 2, \dots$
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 $n = 1, 2, \dots$
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 $S_n(p_n) = 0$.

The game between Forecaster and Reality

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

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

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

Constructing a portmanteau test

In practice, we want to test

1. **calibration (x=1 happens 30% of the times you say p=.3)**
2. **resolution (also true just for times when it rained yesterday)**

For simplicity, consider only calibration.

1. **Use law of large numbers to test calibration for each p.**
2. **Merge the tests for different p.**

FOR $n = 1, 2, \ldots$

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.

Defensive forecasting is not Bayesian

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 knew that a probability can be estimated from a random sample. But this depends on the idd assumption.

Defensive forecasting tells us something new.

1. Our opponent is Reality rather than Nature.
   (Nature follows laws; Reality plays as he pleases.)

2. Defensive forecasting gives probabilities that pass statistical tests regardless of how Reality behaves.

3. I conclude that the idea of an unknown inhomogeneous stochastic process has no empirical content.

Hilary Putnam's counterexample

Hilary Putnam, born 1926, on the right, with Bruno Latour

Why Hilary Putnam thought good probability prediction is impossible. . .

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

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

Reality makes Forecaster look as bad as possible:

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

Skeptic then makes steady money:

$$\text{bet} = \begin{cases} \text{50 cents on 1} & \text{if } p_n < 0.5 \\ \text{50 cents on 0} & \text{if } p_n \geq 0.5 \end{cases}$$

But the example is artificial, because the testing strategy is discontinuous in the forecast p_n.

But Skeptic's move

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

is discontinuous in p. This infinitely abrupt shift—an artificial idealization—is crucial to the counterexample.

Forecaster can defeat any strategy for Skeptic if

- the strategy for Skeptic is continuous in p, or
- Forecaster is allowed to randomize, announcing a probability distribution for p rather than a sharp value for p.

See Working Papers 7 & 8 at www.probabilityandfinance.com.

The game-theoretic framework for probability

1. Pascal’s game-theoretic probability
2. Making Pascal’s game precise
3. Game-theoretic testing and Cournot’s principle
4. Defensive forecasting
5. Is everything stochastic?

Is everything stochastic?

Does every event have an objective probability?

- Andrei Kolmogorov said no.
- Karl Popper said yes.
- I will say yes.

Of course, each response gives a different meaning to the question.

Does every event have an objective probability?

Kolmogorov said **NO.**

Not every event has a definite probability. The assumption that a definite probability in fact exists for a given event under given conditions is a *hypothesis* which must be verified or justified in each individual case.

Great Soviet Encyclopedia, 1951
(quotation abridged)

Andrei Kolmogorov (1903-1987)

Does every event have an objective probability?

Popper said **YES.**

I suggest *a new physical hypothesis*: every experimental arrangement generates propensities which can sometimes be tested by frequencies.

Realism and the Aim of Science, 1983
(quotation abridged)

Karl Popper (1902-1994)

Karl Popper

1. Published *Logik der Forschung* in Vienna in 1935. Translated into English in 1959.
2. Sought a position in Britain, then left Vienna definitively for New Zealand in 1937.
3. Finally obtained a position in Britain in 1946, after becoming celebrated for *The Open Society*.
4. Wrote his lengthy *Postscript* to the *Logik der Forschung* in the 1950s. It was published in three volumes in 1982-1983.

The *Postscript* was published as three books:

1. *Realism and the Aim of Science.* A philosophical foundation for Kolmogorov's measure-theoretic framework for probability.
   My evaluation: Flawed and ill-informed. But important, because the notion of propensities is extremely popular.
2. *The Open Universe: An Argument for Indeterminism.*
   My evaluation: effective and insufficiently appreciated.
3. *Quantum Mechanics and the Schism in Physics*.

Does every event have an objective probability?

- Kolmogorov considered repeatable conditions. He thought the frequency might not be stable.

I agree.

- Popper <u>imagined</u> repetitions. He asserted the existence of a stable "virtual" frequency even if the imagined repetition is impossible.

A major blunder, Most probabilists, statisticians, and econometricians make the same blunder.

- I assume only that the event is embedded in a sequence of events. We can successively assign probabilities that will pass all statistical tests.

Success in online prediction does not demonstrate knowledge of reality. The statistician's skill resides in the choice of the sequence and the kernel, not in modeling.

Giving probabilities for successive events.

Think “stochastic process, unknown probabilities”, not “iid”.

Can I assign probabilities that will pass statistical tests?

1. If you insist that I announce all probabilities before seeing any outcomes, NO.

2. If you always let me see the preceding outcomes before I announce the next probability, then YES.

Three settings for probability

1. **Causal theory.** I must give probabilities for the whole sequence $x_1, x_2, \ldots$ at the outset. (I do not observe $x_1, \ldots, x_{n-1}$ before predicting x_n.) I can succeed only if I have a valid theory. The valid theory may give only upper and lower probabilities, as in Shafer & Vovk (2001).

2. **On-line prediction.** For each n, I must predict x_n after observing $x_1, \ldots, x_{n-1}$. Using defensive forecasting, I can succeed without knowing anything about Reality. (This is why I say that the idea of an unknown inhomogeneous stochastic process has no empirical content). My predictions will be additive probabilities, not merely upper and lower probabilities.

3. **Probability judgement.** I must predict x without having specified a sequence $x_1, \ldots, x_{n-1}$ preceding it. Perhaps I face advocates with different choices for $x_1, \ldots, x_{n-1}$. It may be difficult to provide even upper and lower probabilities. (Dempster-Shafer is one method for this situation.)