

11th IPMU International Conference

The Game-Theoretic Framework for Probability

Glenn Shafer

July 5, 2006

Part I. A new mathematical foundation for probability theory.
Game theory replaces measure theory.

Part II. Application to statistics: Defensive forecasting.
Good probability forecasting is possible.

Part I. A new mathematical foundation for probability theory.

Game theory replaces measure theory.

- **Mathematics:** Classical probability theorems become theorems in game theory (someone has a winning strategy).
- **Philosophy:** Cournot's principle (an event of small probability does not happen) becomes game-theoretic (you do not get rich without risking bankruptcy).

Part II. Application to statistics: Defensive forecasting.

Good probability forecasting is possible.

- We call it **defensive forecasting** because it defends against a portmanteau (quasi-universal) test.
- Your probability forecasts will pass this portmanteau test **even if reality plays against you.**

Defensive forecasting is a radically new method, not encountered in classical or measure-theoretic probability.

Part I. Basics of Game-Theoretic Probability


1. **Pascal & Ville.** Pascal assumed no arbitrage (you cannot make money for sure) in a sequential game. Ville added Cournot's principle (you will not get rich without risking bankruptcy).
2. The strong law of large numbers
3. The weak law of large numbers



Blaise Pascal (1623–1662), as imagined in the 19th century by Hippolyte Flandrin.


Pascal: Fair division

Peter and Paul play for \$100. Paul is behind. Paul needs 2 points to win, and Peter needs only 1.




If the game must be broken off, how much of the \$100 should Paul get?

It is fair for Paul to pay $\$a$ in order to get $\$2a$ if he defeats Peter and $\$0$ if he loses to Peter.



So Paul should get $\$25$.



Modern formulation: If the game on the left is available, the prices above are forced by the principle of no arbitrage.

Binary probability game.

(Here \mathcal{K}_n is Skeptic's capital and s_n is the total stakes.)

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

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

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

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

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

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

No Arbitrage: If Forecaster announces a strategy in advance, the strategy must obey the rules of probability to keep Skeptic from making money for sure.

In other words, the p_n should be conditional probabilities from some probability distribution for y_1, y_2, \dots .

Blaise Pascal

Probability is about **fair prices** in a **sequential game**.

Pascal's concept of fairness: **no arbitrage**.

Jean Ville

A second concept of fairness: **you will not get rich without
risking bankruptcy**.



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

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

If you never bet more than you have, you will not get infinitely rich.

As Ville showed, this is equivalent to the principle that events of small probability will not happen. We call both principles **Cournot's principle**.

Binary probability game when Forecaster uses the strategy given by a probability distribution P .

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

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

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

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

$\mathcal{K}_n := \mathcal{K}_{n-1} + s_n(y_n - P\{Y_n = 1 | Y_1 = y_1, \dots, Y_{n-1} = y_{n-1}\})$.

Restriction on Skeptic: Skeptic must choose the s_n so that $\mathcal{K}_n \geq 0$ for all n no matter how Reality moves.

Two sides of fairness in game-theoretic probability.

Pascal Constraint on Forecaster: Don't let Skeptic make money for sure. (No arbitrage.)

Ville Constraint on Skeptic: Do not risk bankruptcy.
(Cournot's principle says he will then not make a lot of money.)

Part I. Basics of Game-Theoretic Probability

1. Pascal & Ville
2. **The strong law of large numbers (Borel)**. The classic version says the proportion of heads converges to $\frac{1}{2}$ except on a set of measure zero. The game-theoretic version says it converges to $\frac{1}{2}$ unless you get infinitely rich.
3. The weak law of large numbers

Fair-coin game. (Skeptic announces the amount M_n he risks losing rather than the total stakes s_n .)

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

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

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

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

$$\mathcal{K}_n := \mathcal{K}_{n-1} + M_n y_n.$$

Skeptic wins if

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

(2) either $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n y_i = 0$ or $\lim_{n \rightarrow \infty} \mathcal{K}_n = \infty$.

Otherwise Reality wins.

Theorem Skeptic has a winning strategy.

Who wins? Skeptic wins if (1) \mathcal{K}_n is never negative and (2) either

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n y_i = 0 \quad \text{or} \quad \lim_{n \rightarrow \infty} \mathcal{K}_n = \infty.$$

So the theorem says that Skeptic has a strategy that (1) does not risk bankruptcy and (2) guarantees that either the average of the y_i converges to 0 or else Skeptic becomes infinitely rich.

Loosely: The average of the y_i converges to 0 unless Skeptic becomes infinitely rich.

The Idea of the Proof

Idea 1 Establish an account for betting on heads. On each round, bet ϵ of the account on heads. Then Reality can keep the account from getting indefinitely large only by eventually holding the cumulative proportion of heads at or below $\frac{1}{2}(1 + \epsilon)$.
It does not matter how little money the account starts with.

Idea 2 Establish infinitely many accounts. Use the k th account to bet on heads with $\epsilon = 1/k$. This forces the cumulative proportion of heads to stay at $1/2$ or below.

Idea 3 Set up similar accounts for betting on tails. This forces Reality to make the proportion converge exactly to one-half.

Definitions

- A *path* is an infinite sequence $y_1y_2\dots$ of moves for Reality.
- An *event* is a set of paths.
- A *situation* is a finite initial sequence of moves for Reality, say $y_1y_2\dots y_n$.
- \square is the *initial situation*, a sequence of length zero.
- When ξ is a path, say $\xi = y_1y_2\dots$, write ξ^n for the situation $y_1y_2\dots y_n$.

Game-theoretic processes and martingales

- A real-valued function on the situations is a *process*.
- A process \mathcal{P} can be used as a strategy for Skeptic: Skeptic buys $\mathcal{P}(y_1 \dots y_{n-1})$ of y_n Skeptic in situation $y_1 \dots y_{n-1}$.
- A strategy for Skeptic, together with a particular initial capital for Skeptic, also defines a process: Skeptic's *capital process* $\mathcal{K}(y_1 \dots y_n)$.
- We also call a capital process for Skeptic a *martingale*.

Notation for Martingales

Skeptic begins with capital 1 in our game, but we can change the rules so he begins with α .

Write $\mathcal{K}^{\mathcal{P}}$ for his capital process when he begins with zero and follows strategy \mathcal{P} : $\mathcal{K}^{\mathcal{P}}(\square) = 0$ and

$$\mathcal{K}^{\mathcal{P}}(y_1 y_2 \dots y_n) := \mathcal{K}^{\mathcal{P}}(y_1 y_2 \dots y_{n-1}) + \mathcal{P}(y_1 y_2 \dots y_{n-1}) y_n.$$

When he starts with α , his capital process is $\alpha + \mathcal{K}^{\mathcal{P}}$.

The capital processes that begin with zero form a linear space, for

$$\beta \mathcal{K}^{\mathcal{P}} = \mathcal{K}^{\beta \mathcal{P}} \quad \text{and} \quad \mathcal{K}^{\mathcal{P}_1} + \mathcal{K}^{\mathcal{P}_2} = \mathcal{K}^{\mathcal{P}_1 + \mathcal{P}_2}.$$

So the martingales also form a linear space.

Convex Combinations of Martingales

If \mathcal{P}_1 and \mathcal{P}_2 are strategies, and $\alpha_1 + \alpha_2 = 1$, then

$$\alpha_1(1 + \mathcal{K}^{\mathcal{P}_1}) + \alpha_2(1 + \mathcal{K}^{\mathcal{P}_2}) = 1 + \mathcal{K}^{\alpha_1\mathcal{P}_1 + \alpha_2\mathcal{P}_2}.$$

—LHS is the convex combination of two martingales that each begin with capital 1.

—RHS is the martingale produced by the same convex combination of strategies, also beginning with capital 1.

Conclusion: In the game where we begin with capital 1, we can obtain a convex combination of $1 + \mathcal{K}^{\mathcal{P}_1}$ and $1 + \mathcal{K}^{\mathcal{P}_2}$ by splitting our capital into two accounts, one with initial capital α_1 and one with initial capital α_2 . Apply $\alpha_1\mathcal{P}_1$ to the first account and $\alpha_2\mathcal{P}_2$ to the second.

Infinite Convex Combinations: Suppose $\mathcal{P}_1, \mathcal{P}_2, \dots$ are strategies and $\alpha_1, \alpha_2, \dots$ are nonnegative real numbers adding to one.

- If $\sum_{k=1}^{\infty} \alpha_k \mathcal{P}_k$ converges, then $\sum_{k=1}^{\infty} \alpha_k \mathcal{K}^{\mathcal{P}_k}$ also converges.
- $\sum_{k=1}^{\infty} \alpha_k \mathcal{K}^{\mathcal{P}_k}$ is the capital process from $\sum_{k=1}^{\infty} \alpha_k \mathcal{P}_k$.
- You can prove this by induction on

$$\mathcal{K}^{\mathcal{P}}(y_1 y_2 \dots y_n) := \mathcal{K}^{\mathcal{P}}(y_1 y_2 \dots y_{n-1}) + \mathcal{P}(y_1 y_2 \dots y_{n-1}) y_n.$$

In game-theoretic probability, you can usually get an infinite convex combination of martingales, but you have to check on the convergence of the infinite convex combination of strategies. In a sense, this explains the historical confusion about countable additivity in measure-theoretic probability (see Working Paper #4).

The greater power of game-theoretic probability

Instead of a probability distribution for y_1, y_2, \dots , maybe you have only a few prices. Instead of giving them at the outset, maybe you make them up as you go along. Instead of

Skeptic announces $M_n \in \mathbb{R}$.
Reality announces $y_n \in \{-1, 1\}$.
 $\mathcal{K}_n := \mathcal{K}_{n-1} + M_n y_n$.

use

Skeptic announces $M_n \in \mathbb{R}$.
Reality announces $y_n \in [-1, 1]$.
 $\mathcal{K}_n := \mathcal{K}_{n-1} + M_n y_n$.

or

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

Part I. Basics of Game-Theoretic Probability

1. Pascal & Ville
2. The strong law of large numbers. Infinite and impractical: You will not get infinitely rich in an infinite number of trials.
3. **The weak law of large numbers.** Finite and practical: You will not multiply your capital by a large factor in N trials.

The weak law of large numbers (Bernoulli)

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

FOR $n = 1, \dots, N$:

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

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

$\mathcal{K}_n := \mathcal{K}_{n-1} + M_n y_n.$

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

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

Part II. Defensive Forecasting

1. **Thesis.** Good probability forecasting is possible.
2. **Theorem.** Forecaster can beat any test.
3. **Research agenda.** Use proof to translate tests of Forecaster into forecasting strategies.
4. **Example.** Forecasting using LLN (law of large numbers).

THESIS

Good probability forecasting is possible.

We can always give probabilities with good calibration and resolution.

PERFECT INFORMATION PROTOCOL

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*.

FOR $n = 1, 2, \dots$

Reality announces $x_n \in \mathbf{X}$.

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

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

Forecaster can give p s with good calibration and resolution *no matter what Reality does*.

Philosophical implications:

- To a good approximation, everything is stochastic.
- Getting the probabilities right means describing the past well, not having insight into the future.

THEOREM. Forecaster can beat any test.

FOR $n = 1, 2, \dots$

Reality announces $x_n \in \mathbf{X}$.

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

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

- **Theorem.** Given a test, Forecaster has a strategy guaranteed to pass it.
- **Thesis.** There is a test of Forecaster universal enough that passing it implies the p s have good calibration and resolution. (Not a theorem, because “good calibration and resolution” is fuzzy.)

The probabilities are tested by another player, Skeptic.

FOR $n = 1, 2, \dots$

Reality announces $x_n \in \mathbf{X}$.

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

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

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

Skeptic's profit $:= s_n(y_n - p_n)$.

A **test of Forecaster** is a strategy for Skeptic that is continuous in the p s. **If Skeptic does not make too much money, the p s pass the test.**

Theorem If Skeptic plays a known continuous strategy, Forecaster has a strategy guaranteeing that Skeptic never makes money.

This concept of test generalizes the standard stochastic concept.

Stochastic setting:

- There is a probability distribution P for the x s and y s.
- Forecaster uses P 's conditional probabilities as his p s.
- Reality chooses her x s and y s from P .

Standard concept of statistical test:

- Choose an event A whose probability under P is small.
- Reject P if A happens.

In 1939, Jean Ville showed that in the stochastic setting, the standard concept is equivalent to a strategy for Skeptic.

Why insist on continuity? Why count only strategies for Skeptic that are continuous in the p s as tests of Forecaster?

1. *Brouwer's thesis*: A computable function of a real argument is continuous.
2. Classical statistical tests (e.g., reject if LLN fails) correspond to continuous strategies.

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 $S_n : [0, 1] \rightarrow \mathbb{R}$.

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 $n = 1, 2, \dots$

Reality announces $x_n \in \mathbf{X}$.

Skeptic announces continuous $S_n : [0, 1] \rightarrow \mathbb{R}$.

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 $S_n(p_n) = 0$.

Research agenda. Use proof to translate tests of Forecaster into forecasting strategies.

- **Example 1:** Use a strategy for Sceptic that makes money if Reality does not obey the LLN (frequency of $y_n = 1$ overall approximates average of p_n). The derived strategy for Forecaster guarantees the LLN—i.e., its probabilities are calibrated “in the large”.
- **Example 2:** Use a strategy for Skeptic that makes money if Reality does not obey the LLN for rounds where p_n is close to p^* . The derived strategy for Forecaster guarantees calibration for p_n close to p^* .
- **Example 3:** Average the preceding strategies for Skeptic for a grid of values of p^* . The derived strategy for Forecaster guarantees good calibration everywhere.
- **Example 4:** Average over a grid of values of p^* and x^* . Then you get good resolution too.

Example 3: 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.

Example 4: Average over a grid of values of p^* and x^* . (The (p^*, x^*) -strategy makes money if calibration fails for n where (p_n, x_n) is close to (p^*, x^*) .) Then you get good calibration and good resolution.

- Define a metric for $[0, 1] \times \mathbf{X}$ by specifying an inner product space H and a mapping

$$\Phi : [0, 1] \times \mathbf{X} \rightarrow H$$

continuous in its first argument.

- Define a kernel $K : ([0, 1] \times \mathbf{X})^2 \rightarrow \mathbb{R}$ by

$$K((p, x)(p', x')) := \Phi(p, x) \cdot \Phi(p', x').$$

The strategy for Skeptic:

$$S_n(p) := \sum_{i=1}^{n-1} K((p, x_n)(p_i, x_i))(y_i - p_i).$$

Skeptic's strategy:

$$S_n(p) := \sum_{i=1}^{n-1} K((p, x_n)(p_i, x_i))(y_i - p_i).$$

Forecaster's strategy: Choose p_n so that

$$\sum_{i=1}^{n-1} K((p_n, x_n)(p_i, x_i))(y_i - p_i) = 0.$$

The main contribution to the sum comes from i for which (p_i, x_i) is close to (p_n, x_n) . So we need to choose p_n to make (p_n, x_n) close (p_i, x_i) for which $y_i - p_i$ average close to zero.

Choose p_n to make (p_n, x_n) look like (p_i, x_i) for which we already have good calibration/resolution.

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.
- www.glennshafer.com: Most of my published articles, including the two following.
- *Statistical Science*, **21** 70–98, 2006: The sources of Kolmogorov's *Grundbegriffe*.
- *Journal of the Royal Statistical Society, Series B* **67** 747–764, 2005: Good randomized sequential probability forecasting is always possible.

Standard stochastic concept of statistical test:

- Choose an event A whose probability under P is small.
- Reject P if A happens.

Ville's Theorem: In the stochastic setting...

- Given an event of probability less than $1/C$, there is a strategy for Skeptic that turns \$1 into $\$C$ without risking bankruptcy.
- Given a strategy for Skeptic that starts with \$1 and does not risk bankruptcy, the probability that it turns \$1 into $\$C$ or more is no more than $1/C$.

So the concept of a strategy for Skeptic generalizes the concept of testing with events of small probability.

Continuity rules out Dawid's counterexample

FOR $n = 1, 2, \dots$

Skeptic announces continuous $S_n : [0, 1] \rightarrow \mathbb{R}$.

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

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

Skeptic's profit $:= S_n(p_n)(y_n - p_n)$.

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

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

But if Skeptic is forced to approximate S_n by a continuous function of p , then the continuous function will have a zero close to $p = 0.5$, and so Forecaster will set $p_n \approx 0.5$.

THREE APPROACHES TO FORECASTING

FOR $n = 1, 2, \dots$

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 (prediction with expert advice).
2. Start with strategies for **Skeptic**. Improve by averaging (approach of this talk).
3. Start with strategies for **Reality** (probability distributions). Improve by averaging (Bayesian theory).