Laws of probabilities in efficient markets

Vladimir Vovk

Department of Computer Science
Royal Holloway, University of London

Fifth Workshop on Game-Theoretic Probability
and Related Topics
15 November 2014, CIMAT, Guanajuato

What I plan to discuss

In this talk I will:

- Consider two designs of prediction markets (out of three in Jake's tutorial).
- Ask the question: Which markets enforce various laws of probability?

There are few answers.

Simplifying assumption: zero interest rates.

Outline

1. Kinds of markets
   - Traditional markets
   - New market design
2. Laws of probability in traditional markets
3. SLLN in new markets
4. Conclusion

Traditional prediction markets

Standard design: shared with the usual stock markets.

Based on limit orders. Limit orders are put into two queues, and then the market orders are executed instantly.

- limit orders supply liquidity
- market orders consume liquidity

Kinds of markets

Different kinds of traditional markets:

- We are predicting some value which will be settled in the future (prediction markets, such as IEM or Intrade, or futures markets).
- The value is never settled (the usual stock markets); we are "predicting" various future values none of which is definitive.

In both cases we consider a sequence (prediction, outcome, prediction, outcome,...).

In general, the market is predicting a long vector; but for simplicity I will discuss only predicting scalars.

Expectations and probabilities (local)

Suppose at some point the current market value is m and the outcome is x.

- If $x \in \{0, 1\}$, we can interpret m as the market's probability for x.
- If x is not binary (for simplicity, we assume $x \in [-1, 1]$), m is the expectation.

Loss functions and scoring rules

A loss function: $\lambda(m, x)$. Scoring rules are essentially the opposite to loss functions: $-\lambda$.

Examples for $m \in [0, 1]$ and $x \in \{0, 1\}$:

- binary log-loss

$$\lambda(m, x) := \begin{cases} -\log m & \text{if } x = 1 \\ -\log(1 - m) & \text{if } x = 0 \end{cases}$$

- square-loss $\lambda(m, x) = (m - x)^2$

Proper loss functions

- A loss function is proper if, for any $m, m' \in [0,1]$,

$$m\lambda(m,1)+(1-m)\lambda(1-m,0) \le m\lambda(m',1)+(1-m)\lambda(1-m',0)$$

(i.e., it "encourages honesty").
- Binary log-loss and binary square-loss functions are proper.
- The generalized square-loss function $\lambda(m,x) = (m-x)^2$ for $x, m \in [-1,1]$ is also "proper" in the sense that, for any probability measures P on $[-1,1]$ and any $m' \in [-1,1]$,

$$\mathbb{E}((X-m)^2) \le \mathbb{E}((X-m')^2),$$

where $X \sim P$ and $m := \mathbb{E}\, X$.

Market scoring rules

Market scoring rules:

- Trader 0 (the sponsor) announces m_0 and agrees to suffer the loss $\lambda(m_0, x)$.
- Trader k, $k = 1, \dots, K$, announces m_k and agrees to suffer the loss $\lambda(m_k, x)$ in exchange for $\lambda(m_{k-1}, x)$.
- At the moment of settlement (when x becomes known), in addition to what the traders agreed to above, the sponsor gets $\lambda(m_K, x)$.

For every trader (except for the sponsor) making a trade this is profitable on average if they follow their own probability distribution.

The sponsors can lose on average (in a predictable manner).

Outline

1. Kinds of markets
2. Laws of probability in traditional markets
   - Strong law of large numbers (SLLN)
   - Other laws
3. SLLN in new markets
4. Conclusion

SLLN

- This talk: only predicting bounded variables (by, say, 1 in absolute value).
- Simple proofs for traditional and Hanson's (with square loss) markets.

SLLN for bounded random variables for traditional markets

Protocol:

$\mathcal{K}_0 = 1$.
FOR $n = 1, 2, \dots$:
  Forecaster announces $m_n \in \mathbb{R}$.
  Sceptic announces $M_n \in \mathbb{R}$.
  Reality announces $x_n \in [-1, 1]$.
  $\mathcal{K}_n := \mathcal{K}_{n-1} + M_n(x_n - m_n)$.
END FOR.

Skeptic buys tickets paying $x_n - m_n$; $\mathcal{K}_n$: his capital.

Rules of the game

Sceptic wins the game if

- $\mathcal{K}_n$ is never negative
- either

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} (x_i - m_i) = 0$$

or

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

Proposition

Proposition

Sceptic has a winning strategy in this game.

Interpretation: usually based on the principle of the impossibility of a gambling system. Not always; e.g., the predictions in Intrade either allow us to become infinitely rich or are calibrated. Which?

General definition: an event E is almost certain if Sceptic has a strategy that does not risk bankruptcy and makes him infinitely rich if E fails to happen.

Or: Sceptic can force E.

I will almost prove this simple SLLN.

Usual tricks:

- we can replace $\mathcal{K}_n \to \infty$ with $\sup_n \mathcal{K}_n = \infty$ [wait until $\mathcal{K}_n$ reaches C and stop playing; combine this for different $C \to \infty$]
- if $E_1, E_2, \ldots$ are almost certain, $\cap E_i$ is also almost certain [combine the corresponding strategies in the sense of a convex combination; analogous to using σ-additivity]
- suppose $m_n = 0$ for all n

Lemma

Lemma

Suppose $\epsilon > 0$. *Then Sceptic can "weakly force"*

$$\limsup_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} x_i \leq \epsilon.$$

The same argument, with $-\epsilon$ in place of ϵ:

$$\liminf_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} x_i \geq -\epsilon \quad \text{a.s.}$$

Combine this for all ϵ.

Proof of the lemma (1)

Sceptic always buys $\epsilon\mathcal{K}_{n-1}$ at trial n; then

$$\mathcal{K}_n = \prod_{i=1}^{n} (1 + \epsilon x_i) .$$

On the paths where $\mathcal{K}_n$ is bounded:

$$\prod_{i=1}^{n} (1 + \epsilon x_i) \le C; \qquad \sum_{i=1}^{n} \ln (1 + \epsilon x_i) \le D;$$

since $\ln(1 + t) \ge t - t^2$ whenever $t \ge -\frac{1}{2}$,

$$\epsilon \sum_{i=1}^{n} x_i - \epsilon^2 \sum_{i=1}^{n} x_i^2 \le D.$$

Proof of the lemma (2)

$$\epsilon \sum_{i=1}^{n} x_i - \epsilon^2 n \le D;$$

$$\epsilon \sum_{i=1}^{n} x_i \le \epsilon^2 n + D;$$

$$\frac{1}{n} \sum_{i=1}^{n} x_i \le \epsilon + \frac{D}{\epsilon n}.$$

Another strategy

Kumon and Takemura: the strategy of buying

$$\frac{1}{2}\bar{x}_{n-1}\mathcal{K}_{n-1}$$

tickets at trial n (where $\bar{x}_{n-1} := \frac{1}{n-1}\sum_{i=1}^{n-1} x_i$) weakly forces the event

$$\lim_{n\to\infty}\frac{1}{n}\sum_{i=1}^{n} x_i = 0.$$

There are many other laws of probability that have been analyzed for traditional and new markets, including:

- law of the iterated logarithm (only $\leq$ for the given protocol)
- weak law of large numbers
- central limit theorem (one-sided for the given protocol)

But for simplicity I will concentrate on the SLLN.

Outline

SLLN for Hanson's market and square loss

Protocol:

$\mathcal{K}_0 = 0$.
FOR $n = 1, 2, \dots$:
Forecaster announces $m_n \in [-1, 1]$.
Sceptic announces $M_n \in [-1, 1]$.
Reality announces $x_n \in [-1, 1]$.
$\mathcal{K}_n := \mathcal{K}_{n-1} + (x_n - m_n)^2 - (x_n - M_n)^2$.
END FOR.

$\mathcal{K}_n$: Sceptic's capital.

Rules of the SLLN game

Sceptic wins the game if

- either

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} (x_i - m_i) = 0$$

- or

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

There is no condition of bounded debt, but later we will get it almost for free.

Proposition

Proposition

Sceptic has a winning strategy in this game.

General definition: an event E is almost certain if Sceptic has a strategy that makes him infinitely rich if E fails to happen.

Or: Sceptic can force E.

The proof is even simpler than for traditional markets.

A more demanding game

Sceptic wins this game game if

- $\mathcal{K}_n \geq -1$ for all n
- either

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} (x_i - m_i) = 0$$

or

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

Proposition

Sceptic has a winning strategy in this game.

-1 can be replaced by $-\epsilon$, where $\epsilon > 0$ is arbitrarily small.

About the proof

I will also almost prove this SLLN.

The main technical tool: the Aggregating Algorithm (an exponential weights algorithm; could be replaced by other algorithms).

Plays a role analogous to that of Kolmogorov's axiom of σ-additivity.

Aggregating Algorithm (AA)

Proposition

Let $p_1, p_2, \dots$ be non-negative weights summing to 1. The AA (with suitable parameters) defines Learner's strategy in the square-loss (resp. log-loss, resp. binary square-loss) game which guarantees that the following inequality will hold at every trial n and for every Expert$_i$, $i = 1, 2, \dots$,

$$\mathrm{Loss}_n(Learner) \le \mathrm{Loss}_n(Expert_i) + a \ln \frac{1}{p_i},$$

where $a = 2$ (resp. $a = 1$, resp. $a = 1/2$).

Corollary

Therefore, for any sequence of strategies for Sceptic, there exists a strategy ensuring

$$\mathcal{K}_n \geq \mathcal{K}_n^i - a \ln \frac{1}{p_i}.$$

Tricks

Similar tricks:

- we can replace $\mathcal{K}_n \to \infty$ with $\sup_n \mathcal{K}_n = \infty$ [wait until $\mathcal{K}_n$ reaches C and then repeat Forecaster's moves; combine the resulting strategies for different $C \to \infty$]
- if $E_1, E_2, \ldots$ are almost certain, $\cap E_i$ is also almost certain [combine the corresponding strategies]

Proof (1)

The strategy $M_n := m_n + \epsilon$ (truncated if necessary) shows that the following event is almost certain:

$$\exists C \forall n : \sum_{i=1}^{n} (x_t - m_t - \epsilon)^2 - \sum_{i=1}^{n} (x_t - m_t)^2 \geq -C$$

$$\exists C \forall n : 2\epsilon \sum_{i=1}^{n} (x_t - m_t) \leq n\epsilon^2 + C.$$

Proof (2)

Considering separately $\epsilon > 0$ and $\epsilon < 0$:

$$-\epsilon \le \liminf_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n}(x_i - m_i) \le \limsup_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n}(x_i - m_i) \le \epsilon$$

is almost certain. QED

Ensuring that the capital is bounded below

- Mix the strategy with unrestricted capital and the strategy whose capital is 0 (following Forecaster).
- The capital will become bounded below.
- This will sacrifice only a finite amount.
- Taking the 0 strategy with a sufficiently large weight ensures a lower bound arbitrarily close to 0.

A LIL

Proposition

In the previous protocol, Sceptic has a strategy that ensures that

- *either*

$$\limsup_{n\to\infty} \left| \frac{\sum_{i=1}^{n}(x_i - m_i)}{\sqrt{n \ln \ln n}} \right| \leq \frac{1}{\sqrt{2}}$$

- *or*

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

The optimality of this LIL

Proposition

In the previous protocol, Forecaster and Reality have a joint strategy that ensures that

- $$\liminf_{n\to\infty} \left| \frac{\sum_{i=1}^n (x_i - m_i)}{\sqrt{n \ln\ln n}} \right| = -\frac{1}{\sqrt{2}}$$
  $$\limsup_{n\to\infty} \left| \frac{\sum_{i=1}^n (x_i - m_i)}{\sqrt{n \ln\ln n}} \right| = \frac{1}{\sqrt{2}}$$
- *and*
  $$\liminf_{n\to\infty} \mathcal{K}_n < \infty.$$

Connections

Essentially, Sceptic's capital $\mathcal{K}_n$ in the traditional market corresponds to Sceptic's capital $\log \mathcal{K}_n$ in Hanson's market and the log-loss game.

Informally, the capital process is a function $\mathcal{K}$ on the finite sequences $m_1, x_1, \ldots, m_n, x_n$ of Sceptic's opponents such that there exists a strategy for Sceptic leading to capital $\mathcal{K}(m_1, x_1, \ldots, m_n, x_n)$ after the opponents choose $m_1, x_1, \ldots, m_n, x_n$, for all n.

- If $\mathcal{K}$ is a capital process in the traditional market with x_i restricted to $\{0, 1\}$, $\log \mathcal{K}$ will be a capital process in Hanson's market with log-loss game.
- And vice versa.

Outline

1. Kinds of markets
2. Laws of probability in traditional markets
3. SLLN in new markets
4. Conclusion

Conclusion

The space of potential problems is huge, the Cartesian product of the laws of probability and the prediction market designs.

For each law of probability and each market design, can Sceptic force this law of probability for this design?

Perhaps a market design is useful only if Sceptic can force a wide range of laws of probability (it is “proper”)...

Proofs and related results (1)

In the case of traditional markets, all proofs and further information can be found in:

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

Masayuki Kumon and Akimichi Takemura.
On a simple strategy weakly forcing the strong law of large numbers in the bounded forecasting game.
Annals of the Institute of Statistical Mathematics **60**, 801–812, 2008.

Proofs and related results (2)

In the case of Hanson's markets:

Robin Hanson.
Combinatorial information market design.
Information Systems Frontiers **5**, 107–119 (2003).

Vladimir Vovk.
Probability theory for the Brier game.
Theoretical Computer Science (ALT 1997 Special Issue) **261**, 57–79 (2001).

Thank you for your attention!