Gambling under unknown probabilities as a proxy for real world decisions under uncertainty

David Aldous

6 October 2022

Part 1: Introduction and setting

- Background: my "Probability in the Real World" course and web site.
- An Inconvenient Truth: for most interesting future real world events, we don't know a numerical probability.
- So what do we do?

Obviously people have written about this issue in specific contexts, but I never found any introductory discussion suitable for one lecture in an undergraduate course. So here is my "blank slate" attempt. This talk is mathematically elementary, but conceptually a little subtle. Extended write-up at

Gambling under unknown probabilities as a proxy for real world decisions under uncertainty. Amer. Math. Monthly, to appear (with F.T. Bruss).
https://www.stat.berkeley.edu/ aldous/Papers/proxy.pdf

There is an obvious spectrum for future events, with two extremes:

Easy case; where we have lots of relevant data

- Yankees win 2022 World Series (chance 22% from betting odds)
- Democrats control Senate after 2022 elections (chance 53% from prediction markets)

Hard case; little or no relevant data (rare events)

- Larger pandemic (more deaths than from COVID) within 15 years
- Use of tactical nuclear weapons within 1 year

Of course most of the things you or I think about in everyday life are between those extremes.

NATIONAL ACADEMIES *Sciences Engineering Medicine*

NATIONAL ACADEMIES PRESS
Washington, DC

This PDF is available at http://nap.nationalacademies.org/26698

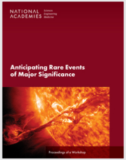

Anticipating Rare Events of Major Significance: Proceedings of a Workshop (2022)

DETAILS

64 pages | 8.5 x 11 | PAPERBACK
ISBN 978-0-309-69306-6 | DOI 10.17226/26698

CONTRIBUTORS

Joe Alper and Liza Hamilton, Rapporteurs; Intelligence Community Studies Board; Division on Engineering and Physical Sciences; National Academies of Sciences, Engineering, and Medicine

BUY THIS BOOK

FIND RELATED TITLES

SUGGESTED CITATION

National Academies of Sciences, Engineering, and Medicine 2022. *Anticipating Rare Events of Major Significance: Proceedings of a Workshop*. Washington, DC: The National Academies Press. https://doi.org/10.17226/26698.

Visit the National Academies Press at nap.edu and login or register to get:

NATIONAL ACADEMIES Sciences Engineering Medicine

Risk Analysis Methods for Nuclear War and Nuclear Terrorism

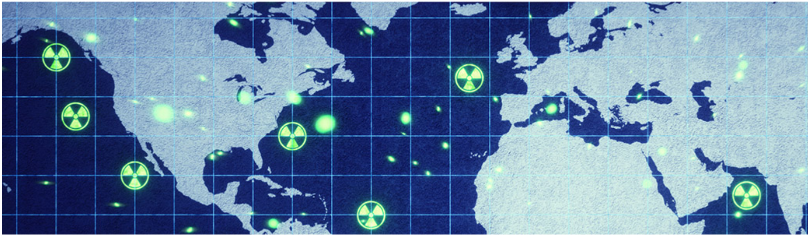

This project will examine the use of analytical methods to assess the risks of nuclear terrorism and nuclear war approaches may play in U.S. security strategy.

Provide feedback on this project

Conceptual aside 1

Each assertion below seems plausible.

- For the 2022 World Series, there is no "true probability" that the Yankees will win.
- Over many bets (e.g. on different sports matches), a gambler who is more accurate at assessing probabilities will do better than a gambler who is less accurate.

Writing this explicitly suggests an inconsistency – do you believe in "true probability", or not?

I adopt the "naive" philosophy that any future event has some unknown true probability. Dogmatic frequentists and dogmatic Bayesians would disagree. But there is solid data (from e.g. prediction tournaments, discussed later) that is difficult to explain within those philosophies.

Conceptual aside 2

Reminder: this talk is envisioning planning a lecture for undergraduates, not showing off fancy research.

Why do we care about probabilities (in the real world)?

There are several general reasons within everyday life – I have a discussion based on data such as 100,000 queries to a search engine containing the strings chance of or probability of – see https://www.cantorsparadise.com/what-is-probability-for-90e5dbc8f7e8 – but the most mathematics-related reason is that we intend to make some **decision** based on a perceived probability and expected utility EU.

My topic in this talk is a small corner of the large "decisions under uncertainty" field, but where the "uncertainty" is in p_{true}. So a throw of a die is not *uncertain*.

Here's the general setting.

- Each individual (or organization) has a "perceived probability" p_{perc} of a given future event.
- Study the "easy cases" where we believe p_{perc} is reasonably close to the unknown p_{true} and is not near 0 or 1.
- How p_{perc} is assessed is very context-dependent; we treat as black box.
- Need to make some decision based on p_{perc}, seeking to maximize EU.

First question: suppose you simply take the action that you would take if $p_{perc} = p_{true}$. What is the cost of error?

cost = EU(if knew p_{true}) - EU(decision based on p_{perc})

Toy example: wedding venue. You are planning a wedding on a given day 9 months ahead (July 6 in Seattle).

- If not raining you prefer venue A (outdoor), if raining you prefer venue B (indoor).
- So you have 4 utilities and want to decide by max EU.
- You obtain p_{perc} (of rain) from historical data for that date (week? month?) for past years (how far back?) – around 15% but surprisingly sensitive.
- Calculate EU(choose A) and EU(choose B) based on p_{perc} and make your decision.

Fact: In this example (exercise) and most such models, the "cost of error" scales as $(p_{perc} - p_{true})^2$ for small errors.
Comment: In continuous examples (how much auto insurance to buy, based on your p_{perc} of accident) this can be viewed as just calculus – we're optimizing a smooth function of p. Not so intuitive in categorical setting.

Fact (repeat): In most such "act as if $p_{perc} = p_{true}$" models, the "cost of error" scales as $(p_{perc} - p_{true})^2$ for small errors.

OK, maybe a rather boring fact. But it offers a different viewpoint on the textbook introduction to sampling. We tell students that the error in estimating a probability from n i.i.d. samples scales as $n^{-1/2}$, but we don't tell them that the **cost** of such errors typically scales as n^{-1}.

There are two more interesting types of questions involving decisions under unknown probabilities. First, recall our "plausible assertion"

- Over many bets (e.g. on different sports matches), a gambler who is more accurate at assessing probabilities will do better than a gambler who is less accurate.

This actually has a precise formalization, next.

Second is the rather vague idea

If your decision involves whether to take a conservative (safe) or aggressive (risky) action, intuition suggests that one should make some allowance for the extra uncertainty (using p_{perc} instead of p_{true}) by being more conservative.

This is challenging, and we don't have any definitive answer.

Part 2: the fundamental example

This basic example can be expressed via two different (mathematically equivalent) stories:

- (first): **A prediction tournament**
- (later): **A Gentleman's Bet**.

The first context is more concrete, with substantial experimental data. The advantage of the second is that it suggests many extensions, for Part 2.

In a **prediction tournament**, contestants state probabilities of future geopolitical events. Here are some questions asked currently on `gjopen.com`.

• Before 1 January 2023, will the presidents of Russia and Ukraine meet in person?

• Before 18 February 2023, will former President Donald Trump be criminally charged with or indicted for a federal and/or state crime in the US?

• Will an electrical blackout lasting at least one hour and affecting 60 million or more people in the US and/or Canada occur before 1 April 2023?

• Before 1 October 2023, will the Israeli Defense Forces (IDF) execute a military strike within the territory of Iran?

• Will the UN declare that a famine exists in any part of a country in the Horn of Africa before 1 April 2023?

• Will there be a lethal confrontation between the national military forces, militia, and/or law enforcement personnel of Taiwan and the People's Republic of China (PRC) before 1 January 2023?

[xxx make prediction on `gjopen`]

How can we assess someone's ability? We do what Gauss said 200 years ago – use **mean square error** MSE. An event is a 0 - 1 variable; if we predict 70% probability then our score on that question is the "squared error":
(if event happens) $(1.0 - 0.70)^2 = 0.09$
(if event doesn't happen) $(0.0 - 0.70)^2 = 0.49$

Your **tournament score** is the sum of scores on each question. As in **golf** one seeks a **low** score. Also as in golf, in a *tournament* all contestants address the same questions; it is not a single-elimination tournament as in tennis.

Write S for your "tournament score" when the true probabilities of the n events are $(p_i, 1 \le i \le n)$ and you predict $(q_i, 1 \le i \le n)$.

- Your **actual** score depends on the (q_i) and the event outcomes.
- Your **expected** score depends on the (q_i) and the (p_i).

A short calculation shows

$$\mathbb{E}S = \sum_i p_i(1-p_i) + n\sigma^2 \quad \text{where} \tag{1}$$

$$\sigma^2 := n^{-1} \sum_i (q_i - p_i)^2$$

σ^2 is your MSE (mean squared error) in **assessing the probabilities**. So for contestants A and B

$$n^{-1}\mathbb{E}(S_A - S_B) = \sigma_A^2 - \sigma_B^2$$

and so **in the long run** we can tell who is the more accurate forecaster **without knowing true probabilities**.

Extensive data, e,g. from IARPA-sponsored prediction tournaments over 2013-2017, shows that some individuals consistently get better scores than others. The natural interpretation is that some individuals are better than others at assessing true probabilities.

What about the short run? The IARPA tournaments had around 100 questions. Under a somewhat realistic model, here are the chances of a more accurate forecaster beating a less accurate forecaster in a 100-question tournament.

		RMS error (less accurate)					
		0.05	0.1	0.15	0.2	0.25	0.3
	0	0.73	0.87	0.95	0.99	1.00	1.00
RMS	0.05		0.77	0.92	0.97	0.99	1.00
error	0.1			0.78	0.92	0.97	0.99
(more	0.15				0.76	0.92	0.97
(accurate)	0.2					0.76	0.91
	0.25						0.73

A gentleman's bet: hypothetical example

You think a future event has probability 20%, your friend thinks it has probability 30%, so you could make a bet at odds corresponding to 25%. For instance

- *you would pay your friend $15 if event did happen*
- *your friend would pay you $5 if event did not happen.*

Each person perceives the bet as favorable.

Gambling odds can be presented in many ways. As in prediction markets (next), we work in term of **contracts**, where one contract on an event will pay $1 if the event occurs, or $0 if not. In this format

You sell 20 contracts to your friend at price 25 (cents) per contract.

[xxx Do real example on `predictit`]

A gentleman's bet: analysis

- Sequence of n future events (arbitrary, unrelated or related).
- Gamblers A and B perceive event i as having probability q_i^A and q_i^B.
- Bet (on event i) at odds corresponding to probability $r_i = (q_i^A + q_i^B)/2$.
- Size of bet (number of contracts) **proportional to difference** $|q_i^A - q_i^B|$.

In our format
A buys $\kappa(q_i^A - q_i^B)$ contracts from B at price r_i per contract.

Here κ is the constant of proportionality – how much money the individual is inclined to gamble with.

Suppose A and B are competing in a prediction tournament (for points) but also betting (for money as above) on each event. Then regardless of outcome, on each event

$$\text{money gain to A from B} \ = \ 2\kappa \ (\text{score of B - score of A}) \ .$$

So the two contexts are mathematically equivalent. But note the specific protocol for gambling.

So our result for prediction tournaments translates to the gambling context: over n bets on different events

$$n^{-1}\mathbb{E}(\text{money gain to A from B}) = \ 2\kappa(\sigma_B^2 - \sigma_A^2)$$

where for each gambler, σ is the RMS error in their probability assessments q_i:

$$\sigma^2 := n^{-1}\sum_i (q_i - p_i)^2.$$

This is true regardless of the unknown true probabilities (p_i).

In a prediction market, the market price represents a consensus probability, so the market acts as Gentleman B. If you follow this protocol with real money in the market, then your long run gain/loss tells you precisely how much better/worse your probability estimates are, compared with the market consensus.

The story so far ......

There are real-world activities in which one can estimate **relative** abilities at estimating probabilities of real-world events.

However .... Prediction tournaments and prediction markets are very special, so maybe not clear how these results might relate to broader "decisions under uncertainty" where probabilities are unknown. In contrast, gambling is a very general activity.

Part 3: explore the general idea

> *In any activity that can be interpreted as of gambling, those agents who are more accurate at estimating probabilities will be more successful than those agents who are less accurate.*

Here we will discuss "toy models", not intended to accurately reflect real world activity for which we have real data. The **bad news** is that the "true probabilities don't matter" property does not extend very far.

In the article we discuss the 5 examples below; today I'll describe the first two..

- **The bookmakers dilemma**: A bookmaker offers odds corresponding to different probabilities, say 64% and 60%, for an event happening or not happening. How to choose these values, based on the bookmakers and the gamblers' perceptions of the probability?
- **Kelly rules**: Adapting to our setting the Kelly criterion for allocating sizes of favorable bets.
- **Bet I'm better than you!**: Two opponents in a game of skill may choose to bet at even odds, but only do so if each believes they are more skillful than the other.

The models above fit into the basic setting where the only unknown quantity is the probability of a given event. The following models have more elaborate settings of "unknowns".

- **Pistols at dawn**: When to fire your one shot, if uncertain about abilities.
- **Unknown consequences of actions**: Unknown mean utilities when choosing or bidding.

General framework

- Make a toy model of a situation where one has to make an action (like deciding whether and how much to bet) whose outcome (gain/loss of money/utility) depends on whether an event of probability p_{true} occurs.
- There is some known optimal (maximize expected utility) action if p_{true} is known.
- But all one has is a "perceived" probability p_{perc}.
- So one just takes the action that one would take if p_{perc} were the true probability.
- Now we study the consequences of the action under the assumption that $p_{perc} = p_{true} + \xi$ for random error ξ, where (for analysis) we usually will assume that ξ has mean zero.

The bookmakers dilemma

A bookmaker offers odds corresponding to different event probabilities, say 64% *and* 60%, *for an event happening or not happening. Here* $[60, 64]$ *is the bookmaker's* **spread**. *How to choose these values, based on the bookmakers' and the gamblers' perceptions of the probability?*

We study some very over-simplified models, to see if the qualitative behavior seems reasonable.

Model 1:. Suppose

- The bookmaker knows the true probability: $p_{book} = p_{true}$.
- Gamblers have different perceived probabilities, say uniform on some interval $[p_{gamb} - L, p_{gamb} + L]$ which is known to the bookmaker.
- The bookmaker offers bets as a spread $[x_1, x_2]$ of contract prices.
- An individual gambler with perceived probability $p_{perc} > x_2$ will buy $\kappa(p_{perc} - x_2)$ contracts at price x_2, and conversely an individual gambler with perceived probability $p_{perc} < x_1$ will sell $\kappa(x_2 - p_{perc})$ contracts at price x_1. Other gamblers do not bet.

In this model, the bookmaker can optimize over x_1 and x_2. Here are the results.

The optimal spread interval is

$$[x_1, x_2] = [\tfrac{2}{3} p_{true} + \tfrac{1}{3}(p_{gamb} - L), \tfrac{2}{3} p_{true} + \tfrac{1}{3}(p_{gamb} + L)] \qquad (2)$$

and the resulting profit is

$$\begin{aligned} \mathbb{E}[\text{ mean gain to bookmaker }] &= \tfrac{2\kappa}{27}(L^2 + 3\Delta^2); \qquad (3) \\ \Delta &:= p_{gamb} - p_{true}. \end{aligned}$$

Recall gambler's perceived probabilities are uniform on the interval $[p_{gamb} - L, p_{gamb} + L]$.

Comments: (a) The bookmaker benefits from the gamblers' "bias" Δ and from the gamblers' spread of perceived probabilities L.
(b) The bookmaker's spread is not centered on p_{true} or on p_{gamb} but on a weighted average.

POLITICS

How Offshore Oddsmakers Made a Killing off Gullible Trump Supporters

The emotions and strategies behind record-setting bets on a MAGA victory that never came.

Model 2: Now suppose the gambler population is unbiased, in the sense that perceived probabilities are uniform on $[p_{true} - L, p_{true} + L]$. But the bookmaker does not know the true probability, but has a perceived probability $p_{book} = p_{true} + \xi$ where the error ξ has a symmetric distribution with variance σ^2. By calculation

The bookmaker's optimal spread interval is symmetric about p_{book} and

$$\mathbb{E}[\text{ mean gain to bookmaker }] = \kappa h(\sigma^2/L^2)L^2$$

for a function h shown in the figure.

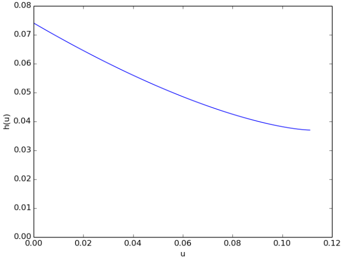

Note the gain depends on the ratio σ/L of error sizes of bookmaker and individual gambler – providing an illustration of the **general theme** that size of errors in perceived probabilities directly affects outcomes.

The Kelly criterion

Consider your stock portfolio.

- For each \$1 in stock i, after 1 year it will be worth $\$X_i$, random.
- Portfolio: invest proportion $q_i \geq 0$ of your money in each stock i.
- So after 1 year your initial fortune is multiplied by $M := \sum_i q_i X_i$.
- The **Kelly optimal** choice of (q_i) is to maximize $\mathbb{E} \log(M)$.

This assumes you know the correct (joint) probability distribution of $(X_i, i \geq 1)$. But no-one does. This fits our "unknown probabilities" setting.

Any realistic analysis would be very complicated. Here is a toy model. Imagine a simple hypothetical setting of betting at **even odds**, on events with probability close to 0.5. If we bet a small proportion a of our fortune and the event occurs with probability $0.5 + \delta$ for small $\delta > 0$ then to first order

$$\text{growth rate } = 2a\delta - a^2/2. \tag{4}$$

So for known $\delta > 0$

- the optimal choice of proportion is $a = 2\delta$
- the resulting optimal growth rate is $2\delta^2$.

Formula (4) remains true for small $\delta < 0$ but of course here the optimal choice is $a = 0$.

In our context there is a perceived probability $0.5 + \delta_{perc}$ and we make the optimal choice based on the perceived probability, that is to bet a proportion $a = \max(0, 2\delta_{perc})$. We use our usual model for perceived probabilities

$$\delta_{perc} = \delta_{true} + \xi.$$

The growth rate $2a\delta_{true} - a^2/2$ can be rewritten as

$$\begin{aligned} \text{growth rate} &= 2(\delta_{true}^2 - \xi^2) \ \text{ if } \xi > -\delta_{true} \\ &= 0 \text{ else.} \end{aligned}$$

Now assume that ξ has Normal$(0, \sigma^2)$ distribution. We can evaluate the expectation of the growth rate in terms of the pdf ϕ and the cdf Φ of the standard Normal Z. For $\delta := \delta_{true}$,

$$\begin{aligned} \mathbb{E}[\text{ growth rate }] &= 2\mathbb{E}[(\delta^2 - \sigma^2 Z^2) 1_{(\sigma Z > -\delta)}] \\ &= 2\left(\delta^2 \Phi(\delta/\sigma) - \sigma^2 S(-\delta/\sigma)\right) \end{aligned}$$

where

$$S(y) := \mathbb{E}[Z^2 1_{(Z>y)}] = y\phi(y) + \Phi(-y).$$

Putting this together,

$$\mathbb{E}[\text{ growth rate }] = 2(\delta_{true}^2 - \sigma^2)\Phi(\delta_{true}/\sigma) + 2\sigma\delta_{true}\phi(\delta_{true}/\sigma). \quad (5)$$

Figure: Growth rate in the Kelly model

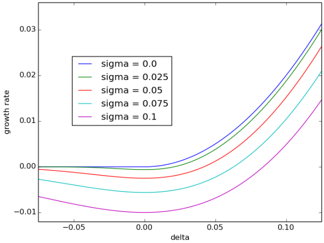

The Figure shows the growth rate as a function of $\delta := \delta_{true}$ for several values of σ. It confirms simple intuition: roughly

if your perceived probabilities have typical error 10% $(\sigma = 0.1)$ *then you will only make money in the long run if the true win probabilities are typically at least* 0.6 $(\delta \geq 0.5 + \sigma)$.

This Kelly setting suggests the general **allowance problem**:
could agents do better if they knew the typical accuracy of their perceived probabilities and adjusted their actions?"

Conclusion

- Project started as material for one lecture in my undergraduate course. Written article on my web page.
- Have not found any comparable introduction to "decisions under unknown probabilities".
- Most frequent question I am asked: isn't this just Bayesian Decision Theory?
- My immediate answer was "we never used Bayes rule".
- Looking at Berger's textbook *Statistical Decision Theory and Bayesian Analysis*, the basic undergrad probability is similar but the focus looks quite different
- We have cuter pictures for our models!

Figure 1: Stories