

UC Berkeley

UC Berkeley Previously Published Works

Title

Gambling Under Unknown Probabilities as Proxy for Real World Decisions Under Uncertainty

Permalink

<https://escholarship.org/uc/item/88j869mt>

Journal

American Mathematical Monthly, 130(4)

ISSN

0002-9890

Authors

Aldous, David J
Bruss, F Thomas

Publication Date

2023-04-21


DOI

10.1080/00029890.2022.2160069

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed



Gambling Under Unknown Probabilities as Proxy for Real World Decisions Under Uncertainty

David J. Aldous & F. Thomas Bruss

To cite this article: David J. Aldous & F. Thomas Bruss (2023): Gambling Under Unknown Probabilities as Proxy for Real World Decisions Under Uncertainty, The American Mathematical Monthly, DOI: [10.1080/00029890.2022.2160069](https://doi.org/10.1080/00029890.2022.2160069)

To link to this article: <https://doi.org/10.1080/00029890.2022.2160069>



Published online: 11 Jan 2023.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)

Gambling Under Unknown Probabilities as Proxy for Real World Decisions Under Uncertainty

David J. Aldous and F. Thomas Bruss

Abstract. The subject of decisions under uncertainty about future events, if lacking sufficient theory or data to make confident probability assessments, poses a challenge for any quantitative analysis. This article suggests one way to first look at this subject. We give elementary examples within a framework for studying decisions under uncertainty where probabilities are only roughly known. The framework, in gambling terms, is that the size of a bet is proportional to the gambler’s perceived advantage based on their perceived probability, and their accuracy in estimating true probabilities is measured by mean-squared-error. Within this framework one can study the cost of estimation errors, and seek to formalize the “obvious” notion that in competitive interactions between agents whose actions depend on their perceived probabilities, those who are more accurate at estimating probabilities will generally be more successful than those who are less accurate.

This article is an extended version of the Brouwer Medal talk at the 2021 Nederlands Mathematisch Congres.

1. INTRODUCTION. The general topic of decisions under uncertainty covers a broad spectrum, from the classical mathematical decision theory surrounding expected utility [19] to modern work (with the celebrated popular exposition by Kahneman [12]) on the cognitive biases exhibited by actual human beings in making such decisions. An important practical point for any discussion is that, for perhaps the majority of real-world decisions we face outside specific professional contexts, numerical probabilities can only be guessed or crudely estimated. Where theory and data do not enable us to estimate probabilities well, can one still seek to study the accuracy of human-estimated probabilities and the effects of errors in such estimation?

Such questions have been discussed at sophisticated levels in various specific contexts (see Section 8 discussion) but this article arises from thinking about how to introduce the topic—what could one say in a single class as part of an undergraduate course in probability?

To pose a starting question, consider the following two intuitive notions.

(A) For unique real-world events, such as “will Netherlands defeat France in their next International Football match,” either there is no “true probability,” or there is some unknown true probability, and we can never tell in any quantitative sense whether your guess of a 30% chance was better than my guess of a 20% chance.

On the other hand, in sports betting or stock market speculation, some individuals do better than others, perhaps more than by pure chance, and one can formulate the following vague general assertion.

(B) In any kind of competitive interactions “under uncertainty” (such as betting on football results) between agents whose actions depend on their perceived probabilities, those who are

doi.org/10.1080/00029890.2022.2160069

MSC: 60G40

more accurate at estimating probabilities will be more successful than those who are less accurate.

In isolation, either (A) or (B) seems very reasonable, but juxtaposed they seem contradictory: do you believe in “true probabilities,” or don’t you? More pedantically, one might say that (A) and (B) represent ends of a spectrum of views about assessment of numerical probabilities for interesting future real-world events.

To the extent that the effects of different outcomes can be expressed within the same quantitative units (as assumed in utility theory), any “decision under uncertainty” model can be regarded as a gambling model. This article introduces the following general framework within which assertion (B) above, interpreted as involving “gambling under unknown probabilities,” can be studied mathematically.

Take 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 usually we need to assume that ξ has mean zero.

As will be discussed in [Section 8](#), this framework is reminiscent of *Bayesian decision theory*, though with a rather different focus.

We will formulate and study three models within this framework, in [Sections 4–6](#). The reader may look ahead to [Section 3](#) for brief descriptions of these models. But let us start in [Section 2](#) with the “fundamental example” of a context in which assertion (B) is verified in a very clear manner.

We emphasize that the mathematics in this article is mostly quite elementary. Our purpose is to show that it is indeed possible to devise quantitative models of various contexts within which (B) is numerically verifiable. So we analyze the behavior of such models and verbally discuss the results of the analyses, rather than proving theorems. What these analyses indicate, in qualitative terms, is that if a gambler is unbiased (ξ has mean zero), then the cost of a small error ξ in perceived probability scales as the variance of ξ . And that in competitive contexts, gamblers with smaller variance will in the long run outperform those with larger variance, providing examples of principle (B). We call this the *squared error principle*. Our models also raise, but do not resolve, another general question we call the *allowance issue*. In our basic setup, gamblers act as if their perceived probabilities were the true probabilities. But if they recognize this is unrealistic, and make some guess about the typical size of their error, should they modify their action to make allowance for their errors? Here our examples do not suggest any general principle, so it remains an intriguing issue for future study.

Further general discussion and pointers to the academic literature are deferred to [Sections 7 and 8](#).

Model ingredients. There are 3 ingredients in most models that we have studied in our framework. First, in gambling-like settings, it is convenient to use the terminology (from prediction markets) of *contracts* rather than *odds*. A contract on an event will pay \$1 if the event occurs, or pay zero if not. In traditional horse race language, one might bet \$7 at odds of 4-to-1 against; that means you would gain \$28 if the horse wins, or lose \$7 if not. This corresponds, in our terminology, to buying 35 contracts at price 0.20 dollars per contract. Here the price is the *implied probability* $p_{implied} = 0.20$ if these were “fair odds.”

Second, we need to model the amount that is bet in any particular case. Of course in most circumstances you would bet on that horse only if your perceived¹ probability p_{perc} is greater than $p_{implied}$. Intuition suggests that you should bet more if your perceived advantage $p_{perc} - p_{implied}$ is larger, and we will use the simplest implementation of that intuition by modeling that the number of contracts bought is proportional to $p_{perc} - p_{implied}$. So the number of contracts bought is $\kappa(p_{perc} - p_{implied})$, where the “constant of proportionality” κ , measuring the scale of an individual’s gambling budget (“affluence,” say), is unimportant for our analyses.

Third, recall our “unknown true probability” p_{true} framework. Any instance of a bet by an individual involves a triple like $(p_{implied}, p_{perc}, p_{true})$ above, and then our model

$$p_{perc} = p_{true} + \xi \quad \text{for random error } \xi$$

will allow us to study mean outcomes in terms of the distribution of the perception error ξ .

This completes a framework for studying the “obvious” vague assertion (B): is it true that, other things being equal, agents with smaller error ξ have better outcomes? Often we will need to make the “unbiased” assumption that the error ξ is such that the expectation $\mathbb{E}[\xi] = 0$ —see Section 7 for discussion of whether this is realistic. Writing $\sigma^2 = \text{var}[\xi]$, the “accuracy of p_{perc} ” can be expressed most usefully as the RMS (root mean square) error σ .

Finally, throughout this article we are **not** envisaging unlikely events with large consequences, for which squared error is clearly not the appropriate measure. In other words, we are implicitly assuming that p_{true} is not close to 0 or 1.

2. THE FUNDAMENTAL EXAMPLE. This example arises in two different, but mathematically equivalent, contexts:

- (first): **A prediction tournament**
- (later): **A gentleman’s bet.**

The first context is more concrete, with substantial experimental data. The advantage of the second context is that it suggests many extensions, which lead to the models in the remaining sections of this article.

A prediction tournament. A *prediction tournament* (see [1] for an introductory account) consists of a collection of questions of the form “state a probability for a specified real-world event happening before a specified date.”² Typical questions in a current (August 2021) tournament at gjopen.com are:

- Before October 31, 2021, will the FDA and/or CDC recommend that at least some Americans fully vaccinated for COVID-19 receive a booster shot?
- Will Haiti hold a presidential election before January 1, 2022?
- What will be the value of the October 2021 S&P/Case-Shiller Home Price Index?
- Before January 1, 2022, will North Korea detonate a nuclear device and/or launch an ICBM with an estimated range of at least 10,000km?
- Will a bitcoin exchange-traded fund trade on a US exchange before January 1, 2022?

¹Terminology is awkward: *estimated* suggests an explicit estimation rule, *subjective* suggests it’s just an opinion. We use *perceived* as a compromise.

²In actual tournaments one can update probabilities as time passes, but for simplicity we consider only a single probability prediction for each question, and only binary outcomes.

- Before January 1, 2022, will the UN declare that a famine exists in Yemen?
- In 2021, will total fire activity in the Amazon exceed the 2020 total count?

These are unique events, not readily analyzed algorithmically. In the popular book [28] contestants are advised to combine what information they can find about the specific event with some “baseline” frequency of roughly analogous previous events. Scoring is by squared error: if you state probability q then on that question

$$\text{your score} = (1 - q)^2 \text{ if event happens; your score} = q^2 \text{ if not.}$$

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.

If you state probability q on a question while the true probability is p , then the expectation of your score on that question, under the true probability, is

$$\begin{aligned} \mathbb{E}[\text{score}] &= p(1 - q)^2 + (1 - p)q^2 \\ &= p(1 - p) + (q - p)^2. \end{aligned}$$

So in a tournament with n questions and unknown true probabilities $(p_i, 1 \leq i \leq n)$, if you state probabilities $(q_i, 1 \leq i \leq n)$ then

$$\mathbb{E}[\text{your tournament score}] = \sum_i p_i(1 - p_i) + \sum_i (q_i - p_i)^2.$$

The first term is the same for all contestants, so if S and \hat{S} are the tournament scores for you and another contestant in an n -question tournament, then

$$n^{-1}\mathbb{E}[S - \hat{S}] = \sigma^2 - \hat{\sigma}^2 \tag{1}$$

where

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

is your RMS error in predicting probabilities and $\hat{\sigma}$ is the other contestant’s RMS error. So we arrive at a key insight: in a prediction tournament, (1) implies that

even though the true probabilities are completely unknown, one can determine (up to small-sample chance variability³) the *relative* abilities of contestants at estimating the true probabilities.

This elementary result seems curiously little-known outside the specific “predict probabilities” community—we give some further discussion in Section 7. A starting point for this article is the observation that it can be reinterpreted as a gambling result, as follows.

³We observe $S - \hat{S}$ rather than its expectation: by a law of large numbers argument, for large n the average $(S - \hat{S})/n$ will be close to its expectation $\mathbb{E}[S - \hat{S}]/n$. See Section 7 for more details.

A gentleman's bet. Suppose two people (A and B) have different perceived probabilities q_A and q_B for a future event and wish to make a bet. Then a contract (to receive 1 dollar if the event occurs) at any price between q_A and q_B is perceived as favorable by each person. Suppose $q_A > q_B$, so A will buy and B will sell, and suppose the price is set at the midpoint $(q_A + q_B)/2 := r$. Recall that the default strategy for A is to buy $\kappa(q_A - r)$ contracts from B.

So A will get some monetary **gain**⁴ from B, via this bet. Now imagine that these two people are also competing in a prediction tournament with the same event as a question. So there will be a **score difference** for this question, which (because A seeks a low score) we write as (score of B) – (score of A).

Consider the relation between the gain and the score difference. Suppose the event occurs. Clearly

$$\text{gain to A} = \kappa(q_A - r)(1 - r)$$

and a brief calculation gives

$$\begin{aligned} \text{score difference} &= (1 - q_B)^2 - (1 - q_A)^2 \\ &= (q_A - q_B)(2 - q_A - q_B) \\ &= 4(q_A - r)(1 - r). \end{aligned}$$

One can check that this relation

$$\text{gain to A} = \frac{\kappa}{4} (\text{score difference})$$

also holds if the event does not happen. So this “gentleman's bet” context is mathematically equivalent to the prediction tournament context. In particular, consider a sequence of n such bets with arbitrary $(q_A(i), q_B(i), p(i))$, $1 \leq i \leq n$, where $p(i)$ is the unknown true probability. Then as at (1) we can write

$$\frac{1}{n} \mathbb{E}[\text{overall gain to A}] = \frac{\kappa}{4} [\sigma_B^2 - \sigma_A^2] \quad (2)$$

where σ_A^2 and σ_B^2 are the mean squared errors of that individual's estimates:

$$\sigma^2 = \frac{1}{n} \sum_i (q(i) - p(i))^2.$$

Again we emphasize the remarkable feature of this result. In the long run, the gambler who is more accurate at perceiving true probabilities will win, with **no assumption about the relation between true and perceived probabilities**.

3. OUR THREE EXAMPLES. We will study three further models of different gambling-like contexts, described briefly below. Alas the “remarkable feature” above is atypical: for other models we will need a probability model of how the perceived and true probabilities are related.

- **The bookmakers dilemma:** A bookmaker offers odds corresponding to different event 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? (Section 4).

⁴A loss is a negative gain.

- **Kelly rules:** Adapting the criterion for allocating sizes of favorable bets. (Section 5).
- **How valuable is it really?:** Unknown utilities when choosing or bidding. (Section 6).

In examining these models, it is worth keeping our two running themes in mind. The *error-squared principle* that we saw in (1, 2) will turn out to hold quite broadly: that is, the quantitative effect of inaccuracy of perceived probabilities scales as σ^2 , the MSE of perceived probabilities. The *allowance issue*—could agents do better by modifying their actions to make allowance for error in perceived probabilities—will remain rather mysterious. We remark that in the fundamental example above, one can show that there is no advantage in misstating your perceived probability, so the allowance issue does not arise. But we suspect that this is another atypical feature of that specific context.

4. THE BOOKMAKERS DILEMMA. This is our most elaborate model, and readers may wish to skip to the subsequent simpler models. We include it here to illustrate possibilities for more sophisticated models within our framework.

Bookmakers are more skilled at predicting the outcomes of games than bettors and systematically exploit bettor biases by choosing prices that deviate from the market clearing price. Levitt [16].

Sports betting is the only [casino] game where you are, in fact, playing against the house You're playing against other people who are actively trying to beat you.⁵ Miller–Davidow [18]

Levitt [16] gives a detailed account of actual bookmaker strategy in the US sports gambling context, and Miller–Davidow [18] describe the nuts and bolts of sports betting from the gambler's viewpoint. A recent example is discussed in [14] under the title *How offshore oddsmakers made a killing off gullible Trump supporters*. Our model below is very crude—let's see whether it is qualitatively reasonable.

In our setup an idealized bookmaker announces a bid price x_1 for a gambler wishing to sell a contract, and an ask price $x_2 > x_1$ for a gambler wishing to buy. That is, in the context on betting whether team T will win an upcoming game, you can bet on “win” by paying x_2 and receiving 1 if T wins, or you can bet on “lose” by paying $1 - x_1$ and receiving 1 if T loses.⁶ In our model, a gambler with perceived probability $p_{perc} > x_2$ will buy $\kappa(p_{perc} - x_2)$ contracts, and a gambler with perceived probability $p_{perc} < x_1$ will sell $\kappa(x_1 - p_{perc})$ contracts. How should the bookmaker choose the *spread interval* $[x_1, x_2]$? If the interval is too wide,⁷ fewer gamblers will bet, whereas if the interval is too narrow then the bookmaker may make too little average profit per bet. Note that we are implicitly assuming a monopoly or cartel of bookmakers, so that gamblers have no alternate venue; otherwise there would be other factors arising from competition between bookmakers.⁸

⁵The point here is that in roulette, blackjack, slot machines etc there is no element of choice on the casino side. In those games, you are not “playing against the house” because the house is not “actively playing,” in the sense of making choices, merely following some prescribed procedure. The only situation where the casino does make a choice is in setting the odds on sports results.

⁶This is equivalent to paying $1 - x_1$ for a contract for the opponent to win.

⁷In US sports betting, the width $x_2 - x_1$ of the spread interval is called the *hold* or the *vig* or the *juice* or the *take* or the *house cut* [18].

⁸Gamblers of course want a small spread interval. Real-world bookmakers present odds in a variety of ways, and (as [18] emphasizes) a prospective gambler should first learn to translate into our *implied odds* format and calculate the spread.

Assume that for a given event the bookmaker knows the distribution (over gamblers) of perceived probabilities, and that this distribution⁹ is uniform over some interval, so we can write this interval as $[p_{gamb} - L, p_{gamb} + L]$, so that p_{gamb} is the consensus probability amongst gamblers. So for a given event¹⁰ the mean gain to bookmaker can be evaluated as

$$\begin{aligned} & \kappa(x_2 - p_{true})\frac{1}{2L} \int_{x_2}^{p_{gamb}+L} (x - x_2)dx + \kappa(p_{true} - x_1)\frac{1}{2L} \int_{p_{gamb}-L}^{x_1} (x_1 - x)dx \\ &= \frac{\kappa}{4L}[(x_2 - p_{true})(p_{gamb} + L - x_2)^2 + (p_{true} - x_1)(x_1 - p_{gamb} + L)^2]. \end{aligned} \quad (3)$$

Note this is assuming

$$[x_1, x_2] \subseteq [p_{gamb} - L, p_{gamb} + L]; \quad (4)$$

which is reasonable because a bookmaker obviously prefers gamblers to have a wide range of perceived probabilities to encourage actual betting, as will be seen in (6, 7).

In the case where the bookmaker knows p_{true} , the bookmaker can optimize (3) over x_1 and x_2 , and the optimal spread interval is

$$[x_1, x_2] = \left[\frac{2}{3}p_{true} + \frac{1}{3}(p_{gamb} - L), \frac{2}{3}p_{true} + \frac{1}{3}(p_{gamb} + L) \right] \quad (5)$$

and the resulting profit is

$$\begin{aligned} \mathbb{E}[\text{mean gain to bookmaker (known } p_{true})] &= \frac{2\kappa}{27}(L^2 + 3\Delta^2); \\ \Delta &:= p_{gamb} - p_{true}. \end{aligned} \quad (6)$$

This is assuming $p_{true} \in [p_{gamb} - L, p_{gamb} + L]$, to satisfy the constraint (4).

In an opposite case where the bookmaker does not know or wish to estimate p_{true} , the bookmaker could instead take the spread interval to be a symmetric interval around p_{gamb} . Applying (3) with $[x_1, x_2] = [p_{gamb} - x, p_{gamb} + x]$ (or calculating directly) we find that the gain is maximized at $x = \frac{1}{3}L$ and the maximized value is

$$\mathbb{E}[\text{mean gain to bookmaker (interval centered at } p_{gamb})] = \frac{2\kappa}{27} L^2. \quad (7)$$

Results (6, 7) are qualitatively what one would expect. Order L^2 arises as the variance of the gamblers' perceived probabilities, and order Δ^2 as the effect of the gambler's bias when the bookmaker is not biased. So the error-squared principle arises again. Note also that the interval in the first case is not symmetric about either p_{true} or p_{gamb} , but rather about a weighted average.

Using the bookmaker's perceived probability. It is unrealistic to assume (as above) that the bookmaker knows the true probability. Instead, can we study a model where, as in the Levitt quote above, the bookmaker is only *more accurate* than the gamblers at assessing probabilities? In our framework we model the bookmaker's perceived probability as

$$p_{book} = p_{true} + \xi$$

⁹One could analyze other distributions.

¹⁰ κ here is the sum of the individual gamblers' affluences κ , assumed independent of their perceived probabilities.

where here the error ξ has a symmetric distribution with variance σ^2 . This leads to our first engagement with the *allowance principle*. Let us continue assuming that the gamblers act as if their perceived probabilities were true probabilities, but let us not require that the bookmaker does so. Note that the bookmaker could do so, that is could use the interval (5) with p_{book} in place of p_{true} . But is this optimal in this model?

In this Section we will calculate the optimal spread interval $[x_1, x_2]$ in terms of (p_{book}, L, σ) under the assumption that $p_{gamb} = p_{true}$.

By a symmetry argument, in the case $p_{gamb} = p_{true}$ we must have $[x_1, x_2] = [p_{book} - y, p_{book} + y]$ for some y . To apply (3),

$$x_2 - p_{true} = y + \xi; \quad p_{gamb} + L - x_2 = L - y - \xi.$$

Neglecting odd orders of ξ (which will have expectation zero)

$$(x_2 - p_{true})(p_{gamb} + L - x_2)^2 = y(L - y)^2 + (y - 2(L - y))\xi^2$$

and so

$$\mathbb{E}[(x_2 - p_{true})(p_{gamb} + L - x_2)^2] = y(L - y)^2 + (3y - 2L)\sigma^2. \quad (8)$$

Maximizing this over y involves setting $d/dy(\cdot) = 0$ and solving the quadratic equation, yielding the optimal value

$$y^* = \frac{1}{3} \left(2 - \sqrt{1 - \frac{9\sigma^2}{L^2}} \right) L$$

provided $\sigma \leq L/3$. So in this model the bookmaker's strategy is to use the spread interval

$$[x_1, x_2] = [p_{book} - y^*, p_{book} + y^*]. \quad (9)$$

In our case where $p_{gamb} = p_{true}$, the value of (8) at y^* works out as

$$\frac{2}{27} \left(1 + \left(1 - \frac{9\sigma^2}{L^2} \right)^{3/2} \right) L^3.$$

We can now use (3) to calculate the mean gain to bookmaker. The contribution from the second term in (3) is the same (by symmetry) as the contribution from the first term, so taking account of the pre-factor $\frac{\kappa}{4L}$ in (3) we find,

$$\mathbb{E}[\text{mean gain to bookmaker}] = \kappa h\left(\frac{\sigma^2}{L^2}\right) L^2 \quad (10)$$

where

$$h(u) = \frac{1}{27} \left(1 + (1 - 9u)^{3/2} \right). \quad (11)$$

The function $h(u)$ is shown in Figure 1 (left). The result corresponds to intuition: as σ increases the gamblers would start to profit from having the correct consensus probability, and so the bookmaker needs to widen the spread interval in response. Note also that $h'(0) < 0$, meaning that the cost of the bookmaker's error scales as σ^2 for small σ , continuing the "error squared" principle.




Figure 1. Mean gain to bookmaker. (Left) The function $h(u)$ at (11) showing gain as a function of bookmaker’s error, if gamblers unbiased. (Right) The function $h^*(u, r)$ at (12) showing gain as a function of gamblers bias, for given bookmaker error.

So the general *allowance issue* question “should agents try to adjust their strategy by making allowance for error in estimating probabilities?” has the answer “yes” for bookmakers in this model, because the optimal spread interval depends on σ . However the practical issue is whether bookmakers could estimate their own σ in order to use a result like (9), leading to deeper issues discussed in Section 7.

Further variations. The assumption above that $p_{gamb} = p_{true}$ is rather pessimistic from the viewpoint of the bookmaker, who presumably will do better in the more realistic case that $\Delta = p_{gamb} - p_{true}$ is non-zero. How to set the spread interval in that case is one of many variants of the models above which may deserve further study. Here we merely note that one can calculate the mean gain to bookmaker, if the $\Delta = 0$ spread interval (9) is used when $\Delta \neq 0$; this gives a formula of the form

$$\mathbb{E}[\text{mean gain to bookmaker}] = \kappa h^*\left(\frac{\sigma^2}{L^2}, \frac{\Delta}{L}\right)L^2. \quad (12)$$

Figure 1 (right) shows the function $r \rightarrow h^*(u, r)$ for $u = 0, 1/36, 2/36$. As r increases, the initial penalty (to bookmaker) of the bookmaker’s inaccuracy is offset increasingly by the inaccuracy of the gamblers’ consensus probability.

5. KELLY RULES. The previous example illustrates the type of “competitive interaction between agents” that our framework is intended for. But for an introductory account, such models seem rather complex, so our remaining two examples are simpler in that they involve only one agent. We can view these as “games against nature”—see Section 7 for discussion.

The general *Kelly strategy* [22], when a range of bets (some favorable) are available, can be illustrated as follows. Suppose that there are 3 available bets, and that bet i will return a random amount $Z_i \geq 0$ from a stake of 1 unit. If the joint distribution of (Z_1, Z_2, Z_3) is known, the strategy is to stake a proportion a_i of one’s fortune on each bet i , where the proportions $\mathbf{a} = (a_1, a_2, a_3)$ are chosen to maximize

$$\mathbb{E}[\log Z_{\mathbf{a}}] := \mathbb{E}[\log(a_1 Z_1 + a_2 Z_2 + a_3 Z_3)]$$

subject to $a_i \geq 0$ and $a_1 + a_2 + a_3 = 1$. So the total return from unit stake will be $Z_{\mathbf{a}}$, and the value $\mathbb{E}[\log Z_{\mathbf{a}}]$ represents the resulting optimal long-term *growth rate*. But what happens when probability distributions are unknown?

In this Section we consider the simple 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 δ , then the first order (we use this approximation throughout)

$$\text{growth rate} = 2a\delta - a^2/2. \tag{13}$$

We view formula (13) as one of the most interesting formulas accessible in introductory mathematical probability. It is curiously hard to find in textbooks, so we repeat the derivation here. For small δ and a ,

$$\begin{aligned} E[\log Z] &= \left(\frac{1}{2} + \delta\right) \log(1 + a) + \left(\frac{1}{2} - \delta\right) \log(1 - a) \\ &\approx \left(\frac{1}{2} + \delta\right)(a - a^2/2) + \left(\frac{1}{2} - \delta\right)(-a - a^2/2) \\ &= 2a\delta - a^2/2. \end{aligned}$$

So for known $\delta > 0$ the optimal choice of proportion is $a = 2\delta$ and the resulting optimal growth rate is $2\delta^2$. Formula (13) remains true for small $\delta < 0$ but of course here the optimal choice is $a = 0$.




Figure 2. Growth rate in the Kelly model, as a function of the true advantage δ .

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 unknown probabilities

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

The growth rate equals $2a\delta_{true} - a^2/2$, which 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 $\delta := \delta_{true}$,

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

where

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

the final equality via an integration by parts. Putting this together,

$$\mathbb{E}[\text{growth rate}] = 2(\delta^2 - \sigma^2)\Phi(\delta/\sigma) + 2\sigma\delta\phi(\delta/\sigma). \quad (14)$$

Figure 2 shows the growth rate as a function of $\delta := \delta_{true}$ for several values of σ . The top curve is the “known advantage $\delta > 0$ ” case, for which the growth rate is $2\delta^2$. In the case where our RMS error in perceived probability is $\sigma > 0$, then even if the game is in fact fair ($\delta = 0$), we will lose money in the long run (growth rate < 0). At first this may seem counter-intuitive, but the Kelley criterion envisages the multiplicative setting: the criterion for growth is $\mathbb{E}[\log Z] > 0$, not $\mathbb{E}[Z] > 1$. It is curious that the “naive” guess, that the criterion for positive growth rate should be “ $\delta > \sigma$ ”, is approximately correct in this model (look where the curves cross the “growth rate = 0” line). This is perhaps merely an artifact of the assumed Normal distribution of error.

The setting here is another setting where the general *allowance issue* question “could agents do better if they knew the typical accuracy of their perceived probabilities and adjusted their actions somewhat?” arises. Intuitively, knowing one’s perception is inaccurate should make one act more conservatively, but in preliminary study we have found it difficult to improve on the growth rate (14). This remains a specific question for future research.

6. HOW VALUABLE IS IT REALLY? The models discussed so far have featured discrete outcomes with unknown probabilities. More generally we can consider continuous outcomes, that is a range of possible utilities, where one’s perception of the utility is inaccurate. Here we study one specific mathematical model chosen to be analytically tractable, and to which two slightly different stories can be attached. Note that in this kind of setting, the details of the model can make a large difference.

Perhaps the simplest type of such model is the “choose the best item from many items” type.

Choosing the best item. Here an agent (you) needs to choose one out of a given set of items. In our model, the true utilities of the items are distributed as the inhomogeneous Poisson process whose intensity function is¹¹

$$\lambda(x) := e^{-x}, \quad -\infty < x < \infty. \quad (15)$$

This distribution arises from classical extreme value theory [21]. In that model the largest true utility $X_{(1)}$ has the Gumbel distribution with c.d.f. and p.d.f.

$$G(x) = \exp(-e^{-x}), \quad g(x) = e^{-x} \exp(-e^{-x}), \quad -\infty < x < \infty.$$

¹¹One imagines there are many items and the distribution of utility score has exponential tail; then re-center the utility score to be near the maximum instead of the mean.

If you knew the true utilities then you would pick that item and gain utility $X_{(1)}$. However, in our model your perceived utility of an item of true utility x is $x + \xi$ for i.i.d. random ξ . So by choosing the item with largest perceived utility, you gain some utility Y which may or may not be $X_{(1)}$, and so you incur a “cost” $X_{(1)} - Y \geq 0$. We will study the mean cost.

Note first that the mean difference between the largest two true utilities, that is $\mathbb{E}[X_{(1)} - X_{(2)}]$, can be calculated easily in terms of

$$N(x) := \text{number of items with true utility } > x$$

because $x \in (X_{(2)}, X_{(1)})$ if and only if $N(x) = 1$. This implies

$$\begin{aligned} \mathbb{E}[X_{(1)} - X_{(2)}] &= \int_{-\infty}^{\infty} \mathbb{P}(N(x) = 1) dx \\ &= \int_{-\infty}^{\infty} e^{-x} \exp(-e^{-x}) dx = 1. \end{aligned} \tag{16}$$

Note also that replacing ξ by $\xi - \mathbb{E}[\xi]$ makes no difference, so we may assume $\mathbb{E}[\xi] = 0$.

Now suppose the errors ξ_σ have $\text{Normal}(0, \sigma^2)$ distribution and write ϕ_σ and Φ_σ for the p.d.f and c.d.f. of that distribution. By basic properties of Poisson processes [15], the process of pairs (true utility of item, perceived utility of item) is Poisson with intensity

$$\lambda_2(x, y) = e^{-x} \phi_\sigma(y - x), \quad -\infty < x, y < \infty.$$

Continuing the calculation of the mean cost could be used as a challenging exercise in an applied probability course; it turns out that the style of argument for (16) can be expanded to show

$$\mathbb{E}[\text{cost}] = \int_{-\infty}^{\infty} \int_{-\infty}^{x_{(1)}} \int_{-\infty}^{\infty} g(x_{(1)}) \Phi_\sigma(y - x_{(1)}) \mathbb{P}(M_{(x, x_{(1)})} < y) \mathbb{P}(M_{(-\infty, x]} \in dy) dx dx_{(1)}$$

where M_I is the maximum perceived utility amongst items with true utility in interval I . And

$$\begin{aligned} \mathbb{P}(M_I < y) &= \exp\left(-\int_I e^{-u}(1 - \Phi_\sigma(y - u)) du\right) \\ \mathbb{P}(M_I \in dy)/dy &= \mathbb{P}(M_I < y) \cdot \int_I e^{-u} \phi_\sigma(y - u) du. \end{aligned}$$

Figure 3 shows numerical values. For small σ we see quadratic growth—the squared error principle again.

Profit or loss at auction? We can reuse the mathematical model at (15) as an auction model, as follows. Now there is one item being sold, and a large number of agents who can bid. For each agent, the future benefit of acquiring the item is some value x which the agent does not know exactly. Instead, the agent perceives the value as $y = x + \xi$ for random ξ , i.i.d. over agents, and the agent bids that perceived value. In this model the points of the Poisson process (15) represent the actual values x_i of the given item to the different agents. The winner of the auction is the agent i whose bid $y_i = x_i + \xi_i$




Figure 3. Mean cost in the “choosing the best item” model, as a function of the RMS error σ in perceived utility.

is largest. In a traditional sealed-bid auction, that winner pays their own bid amount y_i and so makes a profit of $x_i - (x_i + \xi_i) = -\xi_i$. In a modern *Vickrey auction* the winner pays the second-highest bid amount $y_{(2)}$, to mimic a live auction, and so makes a profit of $x_i - y_{(2)}$.




Figure 4. Mean profit for winner in the auction model.

The formula for the mean profit for the winner is complicated, but the model is easily simulated. [Figure 4](#) shows numerical values of the mean profit as a function of σ , when the error ξ has $\text{Normal}(0, \sigma^2)$ distribution. Note that, in the “no error” case $\sigma = 0$, in the traditional auction the actual profit is zero, whereas in the Vickrey auction the mean profit equals 1, by the calculation at (16).

In this model we see the usual “error-squared” quadratic behavior for small σ . It is curious that the difference in mean profits between the two auction protocols remains roughly constant as σ varies.

Game theory aspects. In the auction model above, agents simply bid their perceived value without paying attention to competing bidders. There has been extensive theoretical and applied work on game-theoretic aspects of auctions, in particular in the context of Google auctions for sponsored advertisements [11]. Chapter 14 of [13] gives an introduction to relevant theory. The reader might consider how to combine our framework with game-theoretic models.

7. BACKGROUND AND DISCUSSION.

Remarks on prediction tournaments. In our interpretation of the basic result (1) we are implicitly meaning that *in the long run* we can determine contestants’ relative abilities at estimating probabilities. Of course “long run” assertions deserve some consideration of non-asymptotics. Under fairly plausible more specific assumptions, [1] shows that in a 100-question tournament, if you are 5% more accurate than me (e.g., your RMS error is 15% while mine is 20%), then your chance of beating me is around 75%; increasing to around 92% if a 10% difference in RMS error. So theory says there is quite a lot of chance variability due to event outcomes. However, extensive data, e.g., from IARPA-sponsored prediction tournaments¹² over 2013–2017, shows that some individuals consistently get better scores than others: see [24] for public policy implications. The natural interpretation is that some individuals are better than others at assessing true probabilities. This article has adopted the “naive” philosophy that real-world future events have some unknown true probability. Interpreting the data mentioned above, under alternate philosophical views of probability, strikes us as problematic.

Whether one can estimate σ itself, that is a contestant’s *absolute* rather than *relative* error, is a deep question. Under a certain model, it is shown in [1] that from the distribution of all scores in a 300-player tournament one can determine roughly the σ for the best-scoring contestants, but the assumptions of the model are not readily verifiable.

On unbiased and calibrated estimates. We have not attempted to model how an agent’s perceived probabilities are obtained—just some “black box” method. However, an agent can record predictions and outcomes to check whether the *calibrated* property [27]

the long-run proportion of events with perceived probability near p is approximately p

holds. In principle one can mimic this property in the long run: if one finds that only 15% of events for which one predicted 20% actually occur, then in future when one’s “black box” outputs 20%, one instead predicts 15%. So *calibration* is not so unrealistic for serious gamblers who monitor their performance.

Now imagine a scatter diagram of (p_{perc}, p_{true}) . Freshman statistics reminds us that there are two different regression lines (for predicting one variable given the other). The calibrated property is that one of those lines is the diagonal on $[0, 1]^2$; our *unbiased* property $\mathbb{E}[\xi] = 0$ is that the other line is the diagonal. These are logically different, but if the errors ξ are small then the lines are not very different¹³; so *calibrated* and *unbiased* capture the same intuitive idea.

¹²Contestants recruited via the challenge *can you predict better than the CIA?*

¹³They differ more near 0 or 1, but recall we are implicitly dealing with probabilities not near 0 or 1.

Note also that for simplicity we have usually assumed that the variance of p_{perc} for given p_{true} is a constant σ^2 not depending on p_{true} ; it would be more realistic to relax that assumption.

The error-squared principle in bets against nature. In describing our framework we have emphasized “bets against humans,” such as the gentleman’s bet and the bookmaker’s dilemma, rather than single agent “bets against nature” because, in a sense described below, the latter are simpler. One type of “decisions under unknown probabilities” setting can be abstracted as follows. You have an action which involves choosing a real number x . The outcome depends on whether an event of unknown probability p occurs. The mean gain (in utility) is a function **gain**(x, p). If you knew the true probability p_{true} then you could choose the optimal action

$$x_{opt} = \arg \max_x \mathbf{gain}(x, p_{true}).$$

Instead you have a perceived probability p_{perc} and so you choose an actual action

$$x_{actual} = \arg \max_x \mathbf{gain}(x, p_{perc}).$$

So the mean cost of your error is the difference

$$\mathbf{cost}(p_{perc}) := \mathbf{gain}(x_{opt}, p_{true}) - \mathbf{gain}(x_{actual}, p_{true}).$$

And then calculus tells us that for a generic smooth function **gain** we have $|x_{actual} - x_{opt}|$ is order $|p_{perc} - p_{true}|$ and so because $x \rightarrow \mathbf{gain}(x, p_{true})$ is maximized at x_{opt} ,

$$\mathbf{cost}(p_{perc}) \text{ is order } (p_{perc} - p_{true})^2 \text{ as } p_{perc} \rightarrow p_{true}.$$

So the *cost* of an error in estimating a probability typically scales as the square of the *size* of the error.

In one sense this explains the squared-error principle, but at this level of abstraction it seems hard to appreciate. Note also that one can devise more complicated settings, such as the optimal stopping context in [4], in which one optimizes over discrete strategies x , so that while $p_{true} \rightarrow \mathbf{gain}(x, p_{true})$ is continuous, the optimal strategy for p_{true} switches at some critical value, and this can lead to discontinuity of $p_{perc} \rightarrow \mathbf{cost}(p_{perc})$.

8. REMARKS ON THE ACADEMIC LITERATURE. We do not know any comparable elementary discussion of our general topic—quantitative study of consequences of the fact that numerical probabilities can often only be roughly estimated—but the topic has certainly been considered at a more sophisticated level in many contexts, some of which we describe here.

The most relevant general setting is *Bayesian decision theory* (BDT)—see the well-known textbook [3]. The starting setup in that theory is a loss function $L(\theta, a)$ depending on an unknown parameter θ and a choose-able action a , and a Bayesian would put a prior on the parameter θ and then be able to compute the expectation of loss (as a function of a) and choose the minimizing action a . That theory allows parameters to have arbitrary meaning, whereas in our setting they represent unknown probabilities, so in that sense our basic setting is a very special (and rather atypical) case of BDT. Our “Kelly” example fits directly into the setting of BDT, and indeed an elementary BDT analysis of a similar example is given in [6]. However the central point of this

article is to give a first look at “unknown probabilities” via the conceptual idea (B) and its illustration in the fundamental “gentleman’s bet” example and the more elaborate “bookmaker” example. These examples do not seem to fit naturally into the textbook BDT setting.

Of course one can estimate an unknown probability in the classical context of “repeatable experiments.” In the standard mathematical treatment of stochastic processes models (e.g., [20]), one has a model with several parameters (some relating to probabilities) and the implicit relevance to the real-world is that, with enough data, one can estimate parameters and thereby make predictions. The recent field of *uncertainty quantification* [17,25,26] seeks to address all aspects of error within complicated models, so we are being very simplistic in considering only the uncertainty in the numerical value of a probability. Our focus, exemplified by the prediction tournament examples from Section 2 or sports gambling, is on contexts of “unique events.” That is, cases where one has no causal model or no directly relevant past data to put into a model in order to output a numerical estimate. In such cases (e.g., sports gamblers and bookmakers) there is necessarily some element of human judgment involved in stating a numerical probability.

Sports gambling is too huge a field to survey here; see [18] for a non-mathematical introduction, or browse the *Journal of Quantitative Analysis in Sports*.

Already mentioned is the detailed account of actual bookmaker strategy in the US sports gambling context by Levitt [16]. For one aspect of gambling on horse races to illustrate the style of literature, Green et al. [10] seek to explain the observed low return from betting on longshots in parimutuel markets as follows.

The track deceives naive bettors by suggesting inflated probabilities for longshots and depressed probabilities for favorites. This deception induces naive bettors to underbet favorites, which creates arbitrage opportunities for sophisticated bettors and, from that arbitrage, incremental tax revenue for the track.

At technically sophisticated levels, there has been considerable work on the *calibration* issue—[5, 9] are standard references and [30] provides a recent summary. Closely related is extensive theoretical work on the *forecast aggregation problem*, that is how to combine different probability forecasts into a single consensus forecast. See [2, 8, 23] for representative recent work.

The concept of *Knightian uncertainty*, that is unquantifiable uncertainty, has prompted a wide array of suggested variants of standard mathematical probability. As perhaps the closest topic to our examples, the *Gittins index* [29] is analogous to the Kelly criterion in that it gives an optimal strategy within a certain “multiarmed bandits” context, assuming known distributions. A recent paper [7] describes analysis of the “uncertain distributions” case within a technically sophisticated “nonlinear expectations” framework.

Final remarks. Most readers perhaps do not engage in literal gambling, but a range of more common activities, from fantasy sports to stock market investment, are similar in that decisions involve some implicit judgment of unknown probabilities. Our framework may be conceptually helpful in reminding you to compare your judgment ability with that of the counterparty.

ACKNOWLEDGMENTS. We thank Persi Diaconis, Soumik Pal, Wilfrid Kendall, Venkat Anantharam and two anonymous referees for critiques of early drafts and pointers to technical literature.

REFERENCES

- [1] Aldous, D. J. (2021). A prediction tournament paradox. *Amer. Stat.* 75: 243–248.
- [2] Arieli, I., Y. Babichenko, Y., Smorodinsky, R. (2018). Robust forecast aggregation. *Proc. Natl. Acad. Sci. USA.* 115(52): E12135–E12143.
- [3] Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. New York, NY: Springer-Verlag.
- [4] Bruss, F. T. (2000). Sum the odds to one and stop. *Ann. Probab.* 28(3): 1384–1391.
- [5] Cesa-Bianchi, N., Lugosi, G. (2006). *Prediction, Learning, and Games*. Cambridge, UK: Cambridge Univ. Press.
- [6] Chu, D., Wu, Y., Swartz, T. B. (2019). Modified Kelly criteria. *J. Quant. Anal. Sports* 14: 1–11.
- [7] Cohen, S. N., Treetanthiploet, T. (2021). Gittins’ theorem under uncertainty. arxiv.org/abs/1907.05689
- [8] Cross, D., Ramos, J., Mellers, B., Tetlock, P. E., Scott, D. W. (2018). Robust forecast aggregation: Fourier L_2E regression. *J. Forecast.* 37(3): 259–268.
- [9] Foster, D. P., Vohra, R. V. (1998). Asymptotic calibration. *Biometrika* 85(2): 379–390.
- [10] Green, E., Lee, H., Rothschild, D. M. (2020). The favorite-longshot Midas. Available at: <https://ssrn.com/abstract=3271248>.
- [11] Jansen, B. J., Mullen, T. (2008). Sponsored search: An overview of the concept, history, and technology. *Int. J. Electron. Bus.* 6(2): 114–131.
- [12] Kahneman, D. (2011). *Thinking, Fast and Slow*. New York, NY: Farrar, Straus and Giroux.
- [13] Karlin, A. R., Peres Y. (2017). *Game Theory, Alive*. Providence, RI: American Mathematical Society.
- [14] Kirshner A. (2020). How offshore oddsmakers made a killing off gullible Trump supporters. *Slate*, 17 December. Available at: <https://slate.com/news-and-politics/2020/12/trump-betting-markets-sportsbooks-offshore-2020-election-gambling.html>.
- [15] Günter Last, G., Penrose, M. (2018). *Lectures on the Poisson Process*. Cambridge: Cambridge University Press.
- [16] Levitt, S. D. (2004). Why are gambling markets organised so differently from financial markets? *Econ. J.* 114(495): 223–246.
- [17] McClarren, R. G. (2018). *Uncertainty Quantification and Predictive Computational Science*. New York, NY: Springer-Verlag.
- [18] Miller, E., Davidow, M. (2019). *The Logic of Sports Betting*. Independently published.
- [19] Parmigiani, G., Inoue, L. Y. T. (2009). *Decision Theory*. Chichester, UK: Wiley.
- [20] Pinsky, M. A., Karlin, S. (2011). *An Introduction to Stochastic Modeling*, 4th ed. Amsterdam, Netherlands: Elsevier/Academic Press.
- [21] Resnick, S. I. (1987). *Extreme Values, Regular Variation, and Point Processes*. New York: Springer.
- [22] Rotando, L. M., Thorp, E. O. (1992). The Kelly criterion and the stock market. *Amer. Math. Monthly* 99(10): 922–931.
- [23] Satopää, V. A., Pemantle, R., Ungar, L. H. (2016). Modeling probability forecasts via information diversity. *J. Amer. Stat. Assoc.* 111(516): 1623–1633.
- [24] Scoblic, J. P., Tetlock P. E. (2020). A better crystal ball: The right way to think about the future. *Foreign Affairs*, November/December. Available at: <https://www.foreignaffairs.com/articles/united-states/2020-10-13/better-crystal-ball>.
- [25] Smith, R. C. (2014). *Uncertainty Quantification*. Philadelphia, PA: SIAM
- [26] Soize, C. (2017). *Uncertainty Quantification*. Cham: Springer.
- [27] Tetlock, P. E. (2006). *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton NJ: Princeton Univ. Press.
- [28] Tetlock, P. E., Gardner, D. (2015). *Superforecasting: The Art and Science of Prediction*. New York, NY: Crown.
- [29] Weber, R. (1992). On the Gittins index for multiarmed bandits. *Ann. Appl. Probab.* 2(4): 1024–1033.
- [30] Zhao, S., Ermon, S. (2020). Right decisions from wrong predictions: A mechanism design alternative to individual calibration. arxiv.org/abs/2011.07476

DAVID ALDOUS received his Ph.D. in mathematics from Cambridge University, and has been on the faculty of the U.C. Berkeley Statistics department since 1979. A central research theme has been the study of large finite random structures, obtaining asymptotic behavior as the size tends to infinity via consideration of some suitable infinite random structure. He is currently engaged in articulating critically what mathematical probability says about the real world.

Department of Statistics, University of California at Berkeley
aldousdj@berkeley.edu

F. THOMAS BRUSS obtained his Ph.D. in Mathematics from the University of Saarbrücken in 1977. After an academic career in Germany, UK, Belgium, and USA, in 1993 Thomas obtained the chair “Mathématiques Générales” at U. Libre de Bruxelles (ULB) where he has stayed ever since. Now professor emeritus, Thomas continues research in Applied Probability with main interests in the fields of Branching Processes and Optimal Stopping.

Faculté des Sciences, Université Libre de Bruxelles
tbruss@ulb.ac.be