

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

David J. Aldous and F. Thomas Bruss

April 3, 2021

Abstract

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.

Draft - do not cite

1 Introduction.

The general topic of decisions under uncertainty covers a broad spectrum, from the classical mathematical decision theory surrounding expected utility [12] to modern work (with the celebrated popular exposition by Kahneman [7]) 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 10 discussion) but this article arises from thinking 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 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 more accurate at estimating probabilities will be more successful than those who are less accurate.

Notions (A) and (B) are not exactly contradictory, but 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.

We will formulate and study 5 models within this framework, in sections 4 – 8. 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 very elementary. Further general discussion and pointers to the academic literature are deferred to sections 9 and 10.

1.1 Model ingredients.

There are 3 ingredients in almost all our models. 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, of 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 modelling that the number of contracts bought is proportional to $p_{perc} - p_{implied}$. This is roughly the *Kelly strategy* (section 6). 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

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

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 $\mathbb{E}\xi = 0$ – see section 9.2 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 .

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.

2.1 A prediction tournament.

A *prediction tournament* (see [2] for an introductory account) consists of a collection of questions of the form “state a probability for a specified real-world

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

event happening before a specified date”.² Typical questions in a current (March 2021) tournament at gjopen.com are: before 1 October 2021:

- Will Boris Johnson cease to be prime minister of the United Kingdom?
- Will China officially declare an air-defense identification zone over any part of the South China Sea?
- Will the U.S. leveraged loan default rate exceed 5.0% ?
- Will Saudi Arabia diplomatically recognize the State of Israel?
- Will the UN declare that a famine exists in any part of Yemen?
- Will Russia conduct a flight test of an RS-28 Sarmat ICBM?

These are unique events, not readily analyzed algorithmically. In the popular book [19] 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

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

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

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

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

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.

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 9. A starting point for this article is the observation that it can be re-interpreted as a gambling result, as follows.

2.2 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

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

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

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

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. See section 9.

⁴A loss is a negative gain.

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 list of examples.

We will study 5 further models of different gambling-like contexts, indicated informally here and illustrated in Figure 1. 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)
- **Bet I’m better than you!:** Two opponents in a game may choose to bet at even odds, but only do so if each believes they are more skillful than the other. (section 5)
- **Kelly rules:** Adapting to our setting the Kelly criterion for allocating sizes of favorable bets. (section 6)

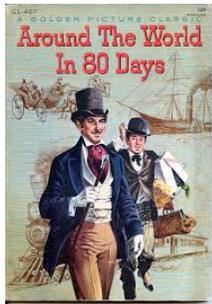
The models above fit into the basic setting where the only unknown quantity is the probability of a given event. The following models have slightly more elaborate settings of “unknowns”.

- **Pistols at dawn:** When to fire your one shot, if uncertain about abilities. (section 7)
- **How valuable is it really?:** Unknown utilities when choosing or bidding. (section 8)

In examining these models, it is worth keeping two themes to keep in mind. The *error-squared principle* that we saw in (1, 2) will hold quite broadly: that is, the quantitative effect of inaccuracy of perceived probabilities scales as σ^2 , the MSE of perceived probabilities. The second theme concerns our default assumption that agents act as if their perceived probability were the true probability – could they instead do better by trying to make allowance for likely error? This *allowance issue* is a fascinating question for future study but we will only address it occasionally in this article.

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. This assumption was also used in identifying “bet proportional to perceived advantage” as approximately the Kelly strategy.

Figure 1: Stories. (Top) gentleman's bet; bookmaker's dilemma; Kelly rules. (Bottom) bet I'm better; pistols at dawn; how valuable.



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* [9].

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. *Miller - Davidow* [11]

Levitt [9] gives a detailed account of actual bookmaker strategy in the U.S. sports gambling context, and Miller - Davidow [11] describes the nuts and bolts of sports betting from the gambler's viewpoint. 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.⁵ 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,⁶ 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.⁷

Assume that for a given event the bookmaker knows the distribution (over gamblers) of perceived probabilities, and that this distribution⁸ is uniform over some interval, so we can write this interval as $[p_{gamb} - L, p_{gamb} + L]$, so that

⁵This is equivalent to paying $1 - x_1$ for a contract for the opponent to win.

⁶In U.S. 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* [11].

⁷Gamblers of course want a small spread interval. Real-world bookmakers present odds in a variety of ways, and (as [11] emphasizes) a prospective gambler should first learn to translate into our *implied odds* format and calculate the spread.

⁸One could analyze other distributions.

p_{gamb} is the consensus probability amongst gamblers. So for a given event⁹

mean gain to bookmaker

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

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) below.

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

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

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. Note the error-squared principle 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.

4.1 Using the bookmaker's perceived probability.

To assume (as above) that the bookmaker knows the true probability is unrealistic. Instead, can we study a model where, as in the Levitt quote above, the

⁹ κ here is the sum of the individual gamblers' affluences κ , assumed independent of their perceived probabilities.

bookmaker is only *more accurate* than the gamblers at assessing probabilities? In our framework we model the bookmaker's perceived probability as

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

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

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

The contribution from the second term in (3) is the same by symmetry, 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

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

The function $h(u)$ is shown in Figure 2 (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.

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 (9), leading to deeper issues discussed in section 9.1.

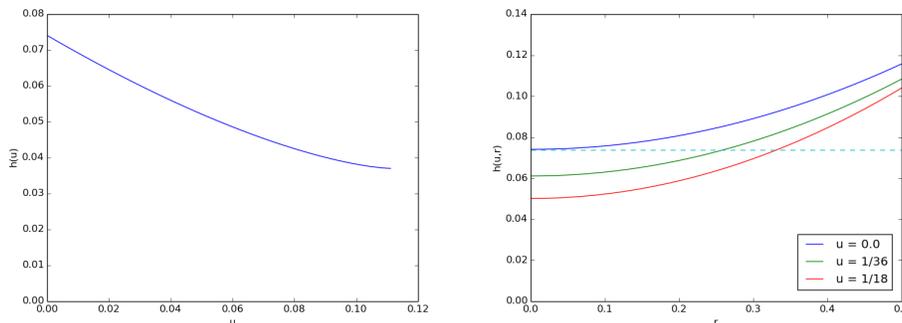
4.2 Further variations.

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 when the $\Delta = 0$ spread interval (9) is used when $\Delta \neq 0$; this gives a formula of the form

$$\mathbb{E}[\text{mean gain to bookmaker}] = \kappa h^*\left(\frac{\sigma^2}{L^2}, \frac{\Delta}{L}\right)L^2. \quad (11)$$

Figure 2 (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.

Figure 2: Mean gain to bookmaker. (Left) as a function of bookmaker's error, if gamblers unbiased. (Right) as a function of gamblers bias, for given bookmaker error.



5 Bet I'm better than you!

The previous examples involve assessing probabilities of events outside your control. A rather different context involves gambling on your own skill. We will give two examples.

5.1 Bradley-Terry games.

Suppose each player has a real-valued “skill” level x , and when player A plays player B

$$\Pr(A \text{ beats } B) = L(x_A - x_B) \quad (12)$$

for a specified function L , the conventional choice being the logistic function

$$L(u) := \frac{e^u}{1 + e^u}, \quad -\infty < u < \infty. \quad (13)$$

This is often called the *Bradley-Terry model* and is closely connected with the *Elo rating* method of estimating skill levels based on past win/lose results – see [1] for an introduction. For our purposes we imagine a game somewhat analogous to a two-person version of poker, in that games only occur when both players agree to play, and all bets are at even odds – you win or lose 1 unit. In this setting, if players knew their skills, then bets would never happen (in our narrow rationality context¹⁰), because the less skillful player would refuse to bet.. But by analogy with the simple 2-person bet in section 2.2, if players don’t know the skills exactly then there will be occasions where each player perceives themselves as better and so are willing to play and bet.

The natural model in our framework is rather complicated, because you may not know your own skill: so there are 4 errors $\xi_{AA}, \xi_{AB}, \xi_{BA}, \xi_{BB}$ to consider, in the format

$$A\text{'s perception of } B\text{'s skill} = x_B + \xi_{AB}.$$

So A is willing to play if $x_A + \xi_{AA} > x_B + \xi_{AB}$, and B is willing to play if $x_B + \xi_{BB} > x_A + \xi_{BA}$. So the game occurs if and only if

$$\xi_{BB} - \xi_{BA} > x_A - x_B > \xi_{AB} - \xi_{AA}.$$

Now suppose that as A varies the quantities $\xi_{AA} - \xi_{AB}$ are independent differently scaled copies of a “standardized” distribution ζ which is symmetric with variance 1; that is

$$\xi_{AA} - \xi_{AB} := \sigma_A \zeta_A, \quad \xi_{BB} - \xi_{BA} := \sigma_B \zeta_B$$

We now have

$$\mathbb{E}(\text{gain to } A) = \Pr(\sigma_A \zeta_A < x_A - x_B) \Pr(\sigma_B \zeta_B > x_A - x_B) (2L(x_A - x_B) - 1). \quad (14)$$

In the context of a given individual A encountering different possible opponents B , we need to average over some distribution for the skill difference $u = x_A - x_B$, and we will take this as uniform on the real line.¹¹ So averaging over B

$$\mathbb{E}(\text{gain to } A) = \int_{-\infty}^{\infty} \Pr(\sigma_A \zeta_A < u) \Pr(\sigma_B \zeta_B > u) (2L(u) - 1) du. \quad (15)$$

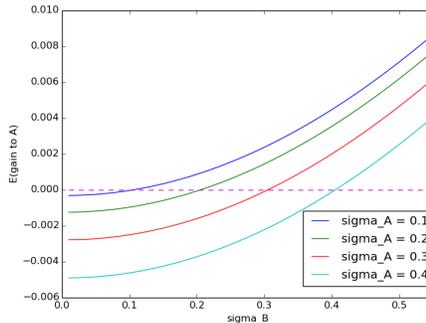
¹⁰In reality there are many reasons one might wish to play against a better player, so it may indeed be rational to accept an unfavorable bet to encourage opponent to play.

¹¹In Bayesian terms this is an improper prior; because bets only occur when when skill difference is small, we get a “proper” distribution of actual bets, although (15) should be more precisely interpreted as a *rate* of gain to A .

Here σ_B may be random but is assumed independent of skill difference u .

We want to determine the effect of the errors in estimating skill – if the variability σ_A of A 's estimate is usually smaller than the variability σ_B of an opponent B 's estimate then we anticipate that A will have positive mean gain. This is not easy to see from the formula (15). If we take σ_B to be constant then Figure 3 shows numerical values. We see the usual error-squared behavior.

Figure 3: Gain to A as a function of opponent error, for given values of A error.



6 Kelly rules.

The general *Kelly strategy* [15], when a range of bets (some favorable) with known probability distributions of outcomes are available, is to divide your fortune (normalized to 1 unit) between bets or reserved, and to do so in the way such that the random value (Z units) of your fortune after the bets are resolved maximizes $\mathbb{E}[\log Z]$. This value $\mathbb{E}[\log Z]$ is the resulting optimal long-term *growth rate*. But what happens when probabilities 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 to first order (we use this approximation throughout)

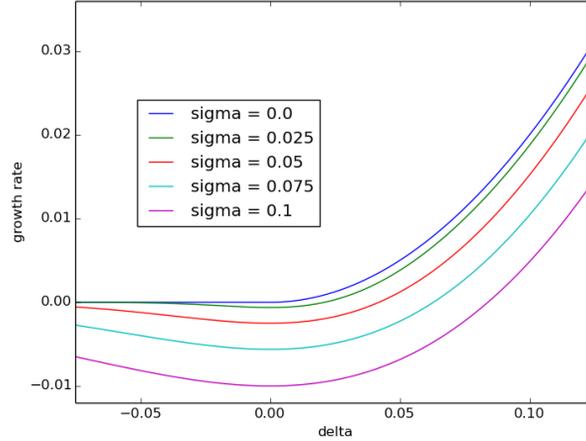
$$\text{growth rate} = 2a\delta - a^2/2. \quad (16)$$

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 (16) 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 unknown probabilities

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

Figure 4: Growth rate in the Kelly model



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 $\text{Normal}(0, \sigma^2)$ 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

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

Putting this together,

$$\mathbb{E}[\text{growth rate}] = 2(\delta^2 - \sigma^2)\Phi(\delta/\sigma) + 2\sigma\delta\phi(\delta/\sigma). \quad (17)$$

Figure 4 shows the growth rate as a function of $\delta := \delta_{true}$ for several values of σ . For $\delta > 0$ we see the usual “quadratic” behavior in both δ and σ .

This 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. In preliminary study we have found it difficult to improve on the growth rate (17).

The “available favorable bets” requirement for Kelly criterion means it is typically used in a stock market context. Unlike our previous examples, there is no interaction with another agent, and we classify this as a “game against nature” as will be discussed in section 9.3.

7 Pistols at dawn.

Let us imagine an 18th century duel: two opponents (A and B) walk toward each other, each with a (not very accurate) pistol with a single shot allowed. As a very simple model, the probability $p(x)$ of a shot incapacitating the opponent depends on the individual and the distance x apart. If the first shot misses, the opponent can then advance and certainly incapacitate the other. Assume A is the worse shot, so the hit probabilities are increasing (as distance x decreases) with $p_B(x) \geq p_A(x)$.

In the case of known probability functions, one can readily see that the natural strategy for A is to shoot (if not previously shot at) at distance around x^* defined as the solution of

$$p_B(x^*) + p_A(x^*) = 1$$

and also the strategy for B is to shoot (if not previously shot at) at the same distance.¹² So

$$\Pr(B \text{ wins}) = p_B(x^*), \Pr(A \text{ wins}) = p_A(x^*). \quad (18)$$

To study the effect of unknown probability functions we need a more detailed model. Assume that each participant has an accuracy parameter $\rho > 1$ and that $p(x) = \min(\rho/x, 1)$. Equation (18) becomes $\rho_A/x + \rho_B/x = 1$ and so $x^* = \rho_A + \rho_B$ and so

$$\Pr(A \text{ wins}) = \rho_A/(\rho_A + \rho_B); \quad \Pr(B \text{ wins}) = \rho_B/(\rho_A + \rho_B). \quad (19)$$

This is in fact equivalent via transformation¹³ to the Bradley-Terry model (12, 13), but in this context participants are compelled to participate rather than having the choice to bet or not to bet. Consider the model where participants know their own ability but not the opponent's. That is,

A perceives B 's accuracy parameter as $\rho_B + \xi_A$;

and B perceives A 's accuracy parameter as $\rho_A + \xi_B$.

So A plans to shoot at distance $x_A^* = \rho_A + \rho_B + \xi_A$; and B plans to shoot at distance $x_B^* = \rho_A + \rho_B + \xi_B$. The outcome of the duel is

- if $\xi_A > \xi_B$ then A shoots first and A wins with probability $\rho_A/(\rho_A + \rho_B + \xi_A)$;
- if $\xi_A < \xi_B$ then B shoots first and so A wins with probability $1 - \rho_B/(\rho_A + \rho_B + \xi_B)$.

Because of the rather arbitrary parametrization, in this model there is a first order effect – if ξ_A and ξ_B have mean zero and the same small variance σ^2 , then the change in probability that A wins scales generically as σ rather than σ^2 . For

¹²At *approximately* the same distance; we can ignore the possibility of simultaneous shots.

¹³Under the re-parametrization $\rho \rightarrow x = \log \rho$. But our error model (additive noise ξ) is different in the two parametrizations.

a more robust observation, if A knows B 's ability but has the opportunity to mis-represent B 's knowledge of A 's ability, which direction of mis-representation is desirable? In this model A would like B to over-estimate A 's ability ($\xi_B > 0$), which would motivate B to shoot earlier and less accurately.

8 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.

8.1 Choosing the best item.

Here one 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 of intensity function

$$\lambda(x) := e^{-x}, \quad -\infty < x < \infty \quad (20)$$

this distribution arising from classical extreme value theory [14]. In that model the largest true utility $X_{(1)}$ has the Gumbel distribution with c.d.f. and p.d.f.

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

If you knew the true utilities then you would pick the 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 can be calculated in terms of

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

via

$$\begin{aligned} \mathbb{E}[X_{(1)} - X_{(2)}] &= \int_{-\infty}^{\infty} \Pr(N(x) = 1) dx & (21) \\ &= \int_{-\infty}^{\infty} e^{-x} \exp(-e^{-x}) dx \\ &= 1. \end{aligned}$$

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. The process of pairs (true utility of item, perceived utility of item) is Poisson with intensity

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

We can calculate the mean cost in the style of (21) as follows.

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

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

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

Figure 5: Mean cost in the “Choosing the best item” model.

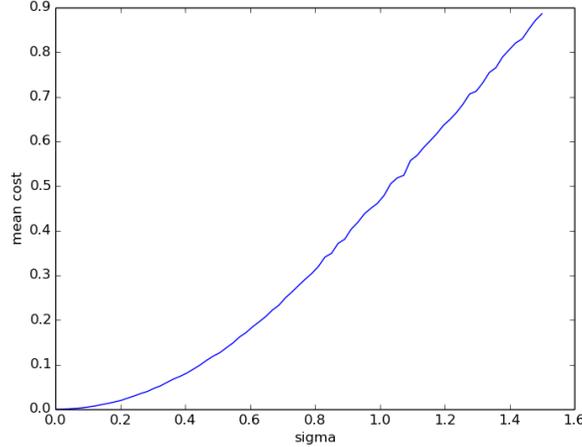


Figure 5 shows numerical values. Note this is what we call a “game against nature” and (as will be discussed in section 9.3) what we find is the expected order σ^2 cost for $\sigma \ll 1$. The subsequent linear growth reflects the model assumption of exponentially growing numbers of items of decreasing utility.

8.2 Profit or loss at auction?

We can re-use the mathematical model at (20) 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, but instead perceives the value as $y = x + \xi$ for random ξ , i.i.d. over agents, and bids that perceived value. In this model the points of the Poisson process (20) 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 $y_i = x_i + \xi_i$ 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 6: Mean gain in the auction model.

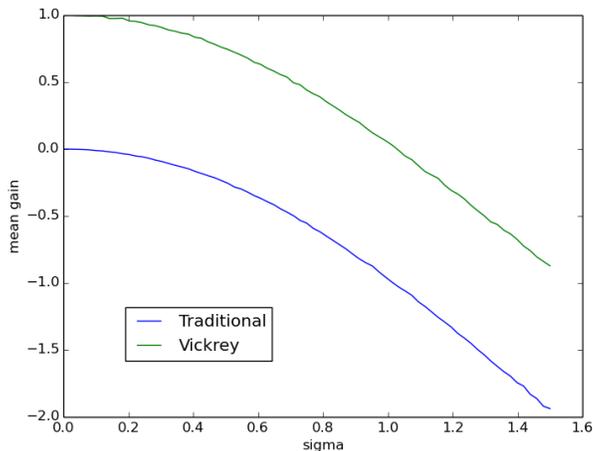


Figure 6 shows numerical values of the mean gain as a function of σ , when ξ has $\text{Normal}(0, \sigma^2)$ distribution. As usual we see σ^2 behavior for small σ . It is curious that the difference in mean gains between the two auction protocols remains roughly constant as σ varies.

Game theory aspects. In the auction model above, agents simply bid their perceived value without 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 [6]. Chapter 14 of [8] gives an introduction to relevant theory. The reader might consider how to combine our framework with game-theoretic models.

9 Background and discussion.

9.1 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, [2] 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 tournaments¹⁴ over 2013-2017, shows that some individuals consistently get better scores than others: see [17] 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 as 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 [2] 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.

9.2 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 [18]

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 different¹⁵; so *calibrated* and *unbiased* capture the same intuitive idea.

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.

¹⁴Contestants recruited via the challenge *can you predict better than the CIA?*

¹⁵They differ more near 0 or 1, but recall we are implicitly dealing with probabilities not near 0 or 1.

9.3 The error-squared principle in bets against nature.

We have emphasized bets against humans because, in a sense described below, “bets against nature” 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 $\mathbf{gain}(x, p)$. 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 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 \mathbf{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 is typically scales as the square of the *size* of the error.

Let’s see how this arises in more detail in one simple example.

Example. Suppose you are planning a wedding, several months ahead, and need to choose now between an outdoor venue A, which you would prefer if you knew it would not be raining, or an indoor venue B, which you would prefer if you knew it would be raining. How to choose? Invoking utility theory (and ignoring the difficulty of actually assigning utilities), there are 4 utilities

- (choose A): utility = a if no rain, utility = b if rain
- (choose B): utility = c if no rain, utility = d if rain

where $a > c$ and $d > b$. Now you calculate the expectation of the utility in terms of the probability p_{true} of rain:

- (choose A): expected utility = $p_{true}b + (1 - p_{true})a$
- (choose B): expected utility = $p_{true}d + (1 - p_{true})c$

There is a critical value p_{crit} where these mean payoffs are equal, and this is the solution of $\frac{p_{crit}}{1-p_{crit}} = \frac{a-c}{d-b}$. Note that p_{crit} depends only on the utilities. If you knew p_{true} your best strategy would be

do action A if $p_{true} < p_{crit}$, do action B if $p_{true} > p_{crit}$.

Instead all we have is our guess p_{perc} , so we use this strategy but based on p_{perc} instead of p_{true} .

What is the cost of not knowing p_{true} ? If p_{true} and p_{perc} are on the same side of p_{crit} then you take the optimal action and there is zero cost; if they are on opposite sides then you take the sub-optimal action and the cost is

$$|p_{true} - p_{crit}|z \text{ where } z = a - b - c + d > 0.$$

Note that z and p_{crit} are determined by the parameter $\theta := (a, b, c, d)$ which we take to be of order 1. In our setting, $(\theta, p_{true}, p_{perc})$ come from some smooth distribution on triples. As in the previous example, the event that $\mathbf{cost} > 0$, and then the mean of \mathbf{cost} given $\mathbf{cost} > 0$, are both order $|p_{perc} - p_{true}|$, and so

$$\mathbb{E}[\mathbf{cost}] \text{ is order } (p_{perc} - p_{true})^2.$$

10 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.

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. [13]) 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* [10] seeks to address all aspects of error within complicated models, but implicitly assuming the model is known to be scientifically accurate or can be calibrated via data. Our focus, exemplified by the prediction tournament examples from section 9.1 or sports gambling, is on contexts of “unique events” where one has no causal models and directly relevant past data, necessitating human judgment not algorithms.

Sports gambling is too huge a field to survey here; see [11] for a non-mathematical introduction, and browse *Journal of Quantitative Analysis in Sports* for articles implicitly relating to events one could bet upon.

Already mentioned is the detailed account of actual bookmaker strategy in the U.S. sports gambling context by Levitt [9]. For one aspect of gambling on horse races to illustrate the style of literature, Green et al. [5] 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.

There has been considerable theoretical work on the *forecast aggregation problem*, that is how to combine different probability forecasts into a single consensus forecast. See [3, 4, 16] for representative recent work. There is also extensive work on the *calibration* issue – see [20] for a recent summary.

References

- [1] David J. Aldous. Elo ratings and the sports model: a neglected topic in applied probability? *Statist. Sci.*, 32(4):616–629, 2017.
- [2] David J. Aldous. A prediction tournament paradox. *The American Statistician*, to appear, 2020.
- [3] I. Arieli, Y. Babichenko, and R. Smorodinsky. Robust forecast aggregation. *Proceedings of the National Academy of Sciences*, 115(52):E12135–E12143, 2018.
- [4] Daniel Cross, Jaime Ramos, Barbara Mellers, Philip E. Tetlock, and David W. Scott. Robust forecast aggregation: Fourier L_2E regression. *J. Forecast.*, 37(3):259–268, 2018.
- [5] Etan Green, Haksoo Lee, and David M. Rothschild. The favorite longshot midas, 2020. Available at <https://ssrn.com/abstract=3271248>.
- [6] Bernard J. Jansen and Tracy Mullen. Sponsored search: an overview of the concept, history, and technology. *International Journal of Electronic Business*, 6(2):114–131, 2008.
- [7] Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.
- [8] Anna R. Karlin and Yuval Peres. *Game theory, alive*. American Mathematical Society, Providence, RI, 2017.
- [9] Steven D. Levitt. Why are Gambling Markets Organised so Differently from Financial Markets? *The Economic Journal*, 114(495):223–246, 04 2004.
- [10] Ryan G. McClarren. *Uncertainty quantification and predictive computational science*. Springer, Cham, 2018.
- [11] Ed Miller and Matthew Davidow. *The Logic Of Sports Betting*. Independently published, 2019.
- [12] Giovanni Parmigiani and Lurdes Y. T. Inoue. *Decision theory*. Wiley Series in Probability and Statistics. John Wiley & Sons, Ltd., Chichester, 2009.
- [13] Mark A. Pinsky and Samuel Karlin. *An introduction to stochastic modeling*. Elsevier/Academic Press, Amsterdam, fourth edition, 2011.
- [14] Sidney I. Resnick. *Extreme Values, Regular Variation, and Point Processes*, volume 4 of *Applied Probability. A Series of the Applied Probability Trust*. Springer-Verlag, New York, 1987.
- [15] Louis M. Rotando and Edward O. Thorp. The Kelly criterion and the stock market. *Amer. Math. Monthly*, 99(10):922–931, 1992.

- [16] Ville A. Satopää, Robin Pemantle, and Lyle H. Ungar. Modeling probability forecasts via information diversity. *J. Amer. Statist. Assoc.*, 111(516):1623–1633, 2016.
- [17] J. Peter Scoblic and Philip E. Tetlock. A better crystal ball: The right way to think about the future. *Foreign Affairs*, November/December, 2020. available at <https://www.foreignaffairs.com/articles/united-states/2020-10-13/better-crystal-ball>.
- [18] Philip E. Tetlock. *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton University Press, 2006.
- [19] Philip E. Tetlock and Dan Gardner. *Superforecasting: The Art and Science of Prediction*. Crown, 2015.
- [20] Shengjia Zhao and Stefano Ermon. Right decisions from wrong predictions: A mechanism design alternative to individual calibration, 2020. arXiv:2011.07476.