
Defensive Forecasting

Vladimir Vovk*
vovk@cs.rhul.ac.uk
http://vovk.net

Akimichi Takemura†
takemura@stat.t.u-tokyo.ac.jp
http://www.e.u-tokyo.ac.jp/~takemura

Glenn Shafer‡*
gshafer@andromeda.rutgers.edu
http://glennshafer.com

Abstract

We consider how to make probability forecasts of binary labels. Our main mathematical result is that for any continuous gambling strategy used for detecting disagreement between the forecasts and the actual labels, there exists a forecasting strategy whose forecasts are ideal as far as this gambling strategy is concerned. A forecasting strategy obtained in this way from a gambling strategy demonstrating a strong law of large numbers is simplified and studied empirically.

1 INTRODUCTION

Probability forecasting can be thought of as a game between two players, Forecaster and Reality:

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

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

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

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

On each round, Forecaster predicts Reality's move y_n chosen from the *label space*, always taken to be $\{0, 1\}$ in this paper. His move, the *probability forecast* p_n , can be interpreted as the probability he attaches to the event $y_n = 1$. To help Forecaster, Reality presents him with an *object* x_n at the beginning of the round; x_n are chosen from an *object space* \mathbf{X} .

Forecaster's goal is to produce p_n that agree with the observed y_n . Various results of probability theory,

*Computer Learning Research Centre, Department of Computer Science, Royal Holloway, University of London, Egham, Surrey TW20 0EX, England.

†Department of Mathematical Informatics, Graduate School of Information Science and Technology, University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan.

‡Rutgers Business School—Newark and New Brunswick, 180 University Avenue, Newark, NJ 07102, USA.

in particular limit theorems (such as the weak and strong laws of large numbers, the law of the iterated logarithm, and the central limit theorem) and large-deviation inequalities (such as Hoeffding's inequality), describe different aspects of agreement between p_n and y_n . For example, according to the strong law of large numbers, we expect that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (y_i - p_i) = 0. \quad (1)$$

Such results will be called *laws of probability* and the existing body of laws of probability will be called *classical probability theory*.

In §2, following [12], we formalize Forecaster's goal by adding a third player, Skeptic, who is allowed to gamble at the odds given by Forecaster's probabilities. We state a result from [14] and [12] suggesting that Skeptic's gambling strategies can be used as tests of agreement between p_n and y_n and that all tests of agreement between p_n and y_n can be expressed as Skeptic's gambling strategies. Therefore, the forecasting protocol with Skeptic provides an alternative way of stating laws of probability.

As demonstrated in [12], many standard proof techniques developed in classical probability theory can be translated into continuous strategies for Skeptic. In §3 we show that for any continuous strategy \mathcal{S} for Skeptic there exists a strategy \mathcal{F} for Forecaster such that \mathcal{S} does not detect any disagreement between the y_n and the p_n produced by \mathcal{F} . This result is a “meta-theorem” that allows one to move from laws of probability to forecasting algorithms: as soon as a law of probability is expressed as a continuous strategy for Skeptic, we have a forecasting algorithm that guarantees that this law will hold; there are no assumptions about Reality, who may play adversarially.

Our meta-theorem is of any interest only if one can find sufficiently interesting laws of probability (expressed as gambling strategies) that can serve as its input. In §4

we apply it to the important properties of unbiasedness in the large and small of the forecasts p_n ((1) is an asymptotic version of the former). The resulting forecasting strategy is automatically unbiased, no matter what data $x_1, y_1, x_2, y_2, \dots$ is observed.

In §5 we simplify the algorithm obtained in §4 and demonstrate its performance on some artificially generated data sets.

2 THE GAMBLING FRAMEWORK FOR TESTING PROBABILITY FORECASTS

Skeptic is allowed to bet at the odds defined by Forecaster's probabilities, and he refutes the probabilities if he multiplies his capital manyfold. This is formalized as a perfect-information game in which Skeptic plays against a team composed of Forecaster and Reality:

BINARY FORECASTING GAME I

Players: Reality, Forecaster, Skeptic

Protocol:

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

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

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

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

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

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

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

Restriction on Skeptic: Skeptic must choose the s_n so that his capital is always nonnegative ($\mathcal{K}_n \geq 0$ for all n) no matter how the other players move.

This is a perfect-information protocol; the players move in the order indicated, and each player sees the other player's moves as they are made. It specifies both an initial value for Skeptic's capital ($\mathcal{K}_0 = 1$) and a lower bound on its subsequent values ($\mathcal{K}_n \geq 0$).

Our interpretation, which will be called the *testing interpretation*, of Binary Forecasting Game I is that \mathcal{K}_n measures the degree to which Skeptic has shown Forecaster to do a bad job of predicting y_i , $i = 1, \dots, n$.

2.1 VALIDITY AND UNIVERSALITY OF THE TESTING INTERPRETATION

As explained in [12], the testing interpretation is valid and universal in an important sense. Let us assume, for simplicity, that objects are absent (formally, that $|\mathbf{X}| = 1$). In the case where Forecaster starts from a probability measure P on $\{0, 1\}^\infty$ and obtains his forecasts $p_n \in [0, 1]$ as conditional probabilities under

P that $y_n = 1$ given y_1, \dots, y_{n-1} , we have a standard way of testing P and, therefore, p_n : choose an event $A \subseteq \{0, 1\}^\infty$ (the *critical region*) with a small $P(A)$ and reject P if A happens. The testing interpretation satisfies the following two properties:

Validity Suppose Skeptic's strategy is measurable and p_n are obtained from P ; \mathcal{K}_n then form a nonnegative martingale w.r. to P . According to Doob's inequality [14, 3], for any positive constant C , $\sup_n \mathcal{K}_n \geq C$ with P -probability at most $1/C$. (If Forecaster is doing a bad job according to the testing interpretation, he is also doing a bad job from the standard point of view.)

Universality According to Ville's theorem ([12], §8.5), for any positive constant ϵ and any event $A \subseteq \{0, 1\}^\infty$ such that $P(A) < \epsilon$, Skeptic has a measurable strategy that ensures $\liminf_{n \rightarrow \infty} \mathcal{K}_n > 1/\epsilon$ whenever A happens, provided p_n are computed from P . (If Forecaster is doing a bad job according to the standard point of view, he is also doing a bad job according to the testing interpretation.) In the case $P(A) = 0$, Skeptic actually has a measurable strategy that ensures $\lim_{n \rightarrow \infty} \mathcal{K}_n = \infty$ on A .

The universality of the gambling scenario of Binary Forecasting Game I is its most important advantage over von Mises's gambling scenario based on subsequence selection; it was discovered by Ville [14].

2.2 CONTINUITY OF GAMBLING STRATEGIES

In [12] we constructed Skeptic's strategies that made him rich when the statement of any of several key laws of probability theory was violated. The constructions were explicit and lead to continuous gambling strategies. We conjecture that every natural result of classical probability theory leads to a continuous strategy for Skeptic.

3 DEFEATING SKEPTIC

In this section we prove the main (albeit very simple) mathematical result of this paper: for any continuous strategy for Skeptic there exists a strategy for Forecaster that does not allow Skeptic's capital to grow, regardless of what Reality is doing. Actually, our result will be even stronger: we will have Skeptic announce his strategy for each round before Forecaster's move on that round rather than making him announce his full strategy at the beginning of the game, and we will drop the restriction on Skeptic. Therefore, we con-

sider the following perfect-information game that pits Forecaster against the two other players:

BINARY FORECASTING GAME II

Players: Reality, Forecaster, Skeptic

Protocol:

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

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

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

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

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

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

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

Theorem 1 *Forecaster has a strategy in Binary Forecasting Game II that ensures $\mathcal{K}_0 \geq \mathcal{K}_1 \geq \mathcal{K}_2 \geq \dots$.*

Proof Forecaster can use the following strategy to ensure $\mathcal{K}_0 \geq \mathcal{K}_1 \geq \dots$:

- if the function $S_n(p)$ takes the value 0, choose p_n so that $S_n(p_n) = 0$;
- if S_n is always positive, take $p_n := 1$;
- if S_n is always negative, take $p_n := 0$. ■

4 EXAMPLES OF GAMBLING STRATEGIES

In this section we discuss strategies for Forecaster obtained by Theorem 1 from different strategies for Skeptic; the former will be called *defensive forecasting strategies*. There are many results of classical probability theory that we could use, but we will concentrate on the simple strategy described in [12], p. 69, for proving the strong law of large numbers.

If $S_n(p) = S_n$ does not depend on p , the strategy from the proof of Theorem 1 makes Forecaster choose

$$p_n := \begin{cases} 0 & \text{if } S_n < 0 \\ 1 & \text{if } S_n > 0 \\ 0 \text{ or } 1 & \text{if } S_n = 0. \end{cases}$$

The basic procedure described in [12] (p. 69) is as follows. Let $\epsilon \in (0, 0.5]$ be a small number (expressing our tolerance to violations of the strong law of large numbers). In Binary Forecasting Game I, Skeptic can ensure that

$$\sup_n \mathcal{K}_n < \infty \implies \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (y_i - p_i) \leq \epsilon \quad (2)$$

using the strategy $s_n = s_n^\epsilon := \epsilon \mathcal{K}_{n-1}$. Indeed, since

$$\mathcal{K}_n = \prod_{i=1}^n (1 + \epsilon(y_i - p_i)),$$

on the paths where \mathcal{K}_n is bounded we have

$$\begin{aligned} \prod_{i=1}^n (1 + \epsilon(y_i - p_i)) &\leq C, \\ \sum_{i=1}^n \ln(1 + \epsilon(y_i - p_i)) &\leq \ln C, \\ \epsilon \sum_{i=1}^n (y_i - p_i) - \epsilon^2 \sum_{i=1}^n (y_i - p_i)^2 &\leq \ln C, \end{aligned}$$

$$\epsilon \sum_{i=1}^n (y_i - p_i) \leq \ln C + \epsilon^2 n,$$

$$\frac{1}{n} \sum_{i=1}^n (y_i - p_i) \leq \frac{\ln C}{\epsilon n} + \epsilon$$

(we have used the fact that $\ln(1+t) \geq t - t^2$ when $|t| \leq 0.5$). If Skeptic wants to ensure

$$\begin{aligned} \sup_n \mathcal{K}_n < \infty &\implies -\epsilon \leq \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (y_i - p_i) \\ &\leq \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (y_i - p_i) \leq \epsilon, \end{aligned}$$

he can use the strategy $s_n := (s_n^\epsilon + s_n^{-\epsilon})/2$, and if he wants to ensure

$$\sup_n \mathcal{K}_n < \infty \implies \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (y_i - p_i) = 0, \quad (3)$$

he can use a convex mixture of $(s_n^\epsilon + s_n^{-\epsilon})/2$ over a sequence of ϵ converging to zero. There are also standard ways of strengthening (3) to

$$\liminf_{n \rightarrow \infty} \mathcal{K}_n < \infty \implies \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (y_i - p_i) = 0;$$

for details, see [12].

In the rest of this section we will draw on the excellent survey [2]. We will see how Forecaster defeats increasingly sophisticated strategies for Skeptic.

4.1 UNBIASEDNESS IN THE LARGE

Following Murphy and Epstein [7], we say that Forecaster is *unbiased in the large* if (1) holds. Let us first consider the one-sided relaxed version of this property

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (y_i - p_i) \leq \epsilon. \quad (4)$$

The strategy for Skeptic described above, $S_n(p) := \epsilon \mathcal{K}_n$, leads to Forecaster always choosing $p_n := 1$; (4) is then satisfied in a trivial way.

Forecaster's strategy corresponding to the two-sided version

$$\begin{aligned} -\epsilon &\leq \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (y_i - p_i) \\ &\leq \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (y_i - p_i) \leq \epsilon \end{aligned} \quad (5)$$

is not much more reasonable. Indeed, it can be represented as follows. The initial capital 1 is split evenly between two accounts, and Skeptic gambles with the two accounts separately. If at the outset of round n the capital on the first account is \mathcal{K}_{n-1}^1 and the capital on the second account is \mathcal{K}_{n-1}^2 , Skeptic plays $s_n^1 := \epsilon \mathcal{K}_{n-1}^1$ with the first account and $s_n^2 := -\epsilon \mathcal{K}_{n-1}^2$ with the second account; his total move is

$$\begin{aligned} S_n(p) &:= \epsilon \mathcal{K}_{n-1}^1 - \epsilon \mathcal{K}_{n-1}^2 \\ &= \epsilon \left(\prod_{i=1}^{n-1} (1 + \epsilon(y_i - p_i)) - \prod_{i=1}^{n-1} (1 + \epsilon(p_i - y_i)) \right). \end{aligned}$$

Therefore, Forecaster's move is $p_n := 1$ if

$$\sum_{i=1}^{n-1} \ln(1 + \epsilon(y_i - p_i)) > \sum_{i=1}^{n-1} \ln(1 + \epsilon(p_i - y_i)),$$

$p_n := 0$ if

$$\sum_{i=1}^{n-1} \ln(1 + \epsilon(y_i - p_i)) < \sum_{i=1}^{n-1} \ln(1 + \epsilon(p_i - y_i)),$$

and p_n can be chosen arbitrarily in the case of equality. The limiting form of this strategy as $\epsilon \rightarrow 0$ is: Forecaster's move is $p_n := 1$ if

$$\sum_{i=1}^{n-1} (y_i - p_i) > 0,$$

$p_n := 0$ if

$$\sum_{i=1}^{n-1} (y_i - p_i) < 0,$$

and p_n can be chosen arbitrarily in the case of equality.

We can see that unbiasedness in the large does not lead to interesting forecasts: Forecaster fulfils his task too well. In the one-sided case (4), he always chooses $p_n := 1$ making

$$\sum_{i=1}^n (y_i - p_i)$$

as small as possible. In the two-sided case (5) with $\epsilon \rightarrow 0$, he manages to guarantee that

$$\left| \sum_{i=1}^n (y_i - p_i) \right| \leq 1. \quad (6)$$

His goals are achieved with categorical forecasts, $p_n \in \{0, 1\}$.

In the rest of this section we consider the more interesting case where $S_n(p)$ depends on p .

4.2 UNBIASEDNESS IN THE SMALL

We now consider a subtler requirement that forecasts should satisfy, which we introduce informally. We say that the forecasts p_n are *unbiased in the small* (or reliable, or valid, or well calibrated) if, for any $p^* \in [0, 1]$,

$$\frac{\sum_{i=1, \dots, n: p_i \approx p^*} y_i}{\sum_{i=1, \dots, n: p_i \approx p^*} 1} \approx p^* \quad (7)$$

provided $\sum_{i=1, \dots, n: p_i \approx p^*} 1$ is not too small.

Let us first consider just one value for p^* . Instead of the ‘‘crisp’’ point p^* we will consider a ‘‘fuzzy point’’ $I : [0, 1] \rightarrow [0, 1]$ such that $I(p^*) = 1$ and $I(p) = 0$ for all p outside a small neighborhood of p^* . A standard choice would be something like $I := \mathbb{I}_{[p_-, p_+]}$, where $[p_-, p_+]$ is a short interval containing p and $\mathbb{I}_{[p_-, p_+]}$ is its indicator function, but we will want I to be continuous (it can, however, be arbitrarily close to $\mathbb{I}_{[p_-, p_+]}$).

The strategy for Skeptic ensuring (2) can be modified as follows. Let $\epsilon \in (0, 0.5]$ be again a small number. Now we consider the strategy $S_n(p) = S_n^{\epsilon, I}(p) := \epsilon I(p) \mathcal{K}_{n-1}$. Since

$$\mathcal{K}_n = \prod_{i=1}^n (1 + \epsilon I(p_i)(y_i - p_i)),$$

on the paths where \mathcal{K}_n is bounded we have

$$\prod_{i=1}^n (1 + \epsilon I(p_i)(y_i - p_i)) \leq C,$$

$$\sum_{i=1}^n \ln(1 + \epsilon I(p_i)(y_i - p_i)) \leq \ln C,$$

$$\epsilon \sum_{i=1}^n I(p_i)(y_i - p_i) - \epsilon^2 \sum_{i=1}^n I^2(p_i)(y_i - p_i)^2 \leq \ln C,$$

$$\epsilon \sum_{i=1}^n I(p_i)(y_i - p_i) - \epsilon^2 \sum_{i=1}^n I(p_i) \leq \ln C$$

(the last step involves replacing $I^2(p_i)$ with $I(p_i)$; the loss of precision is not great if I is close to $\mathbb{I}_{[p_-, p_+]}$),

$$\epsilon \sum_{i=1}^n I(p_i)(y_i - p_i) \leq \ln C + \epsilon^2 \sum_{i=1}^n I(p_i),$$

$$\frac{\sum_{i=1}^n I(p_i)(y_i - p_i)}{\sum_{i=1}^n I(p_i)} \leq \frac{\ln C}{\epsilon \sum_{i=1}^n I(p_i)} + \epsilon.$$

The last inequality shows that the mean of y_i for p_i close to p^* is close to p^* provided we have observed sufficiently many such p_i ; its interpretation is especially simple when I is close to $\mathbb{I}_{[p^-, p^+]}$.

In general, we may consider a mixture of $S_n^{\epsilon, I}(p)$ and $S_n^{-\epsilon, I}(p)$ for different values of ϵ and for different I covering all $p^* \in [0, 1]$. If we make sure that the mixture is continuous (which is always the case for continuous I and finitely many ϵ and I), Theorem 1 provides us with forecasts that are unbiased in the small.

4.3 USING THE OBJECTS

Unbiasedness, even in the small, is only a necessary but far from sufficient condition for good forecasts: for example, a forecaster who ignores the objects x_n can be perfectly calibrated, no matter how much useful information x_n contain. (Cf. the discussion of resolution in [2]; we prefer not to use the term “resolution”, which is too closely connected with the very special way of probability forecasting based on sorting and labeling.) It is easy to make the algorithm of the previous subsection take the objects into account: we can allow the test functions I to depend not only on p but also on the current object x_n ; $S_n(p)$ then becomes a mixture of

$$S_n^{\epsilon, I}(p) := \epsilon I(p, x_n) \prod_{i=1}^{n-1} (1 + \epsilon I(p_i, x_i)(y_i - p_i))$$

and $S_n^{-\epsilon, I}(p)$ (defined analogously) over ϵ and I .

4.4 RELATION TO A STANDARD COUNTER-EXAMPLE

Suppose, for simplicity, that objects are absent ($|\mathbf{X}| = 1$). The standard construction from Dawid [1] showing that no forecasting strategy produces forecasts p_n that are unbiased in the small for all sequences is as follows. Define an infinite sequence y_1, y_2, \dots recursively by

$$y_n := \begin{cases} 1 & \text{if } p_n < 0.5 \\ 0 & \text{otherwise,} \end{cases}$$

where p_n is the forecast produced by the forecasting strategy after seeing y_1, \dots, y_{n-1} . For the forecasts $p_n < 0.5$ we always have $y_n = 1$ and for the forecasts $p_n \geq 0.5$ we always have $y_n = 0$; obviously, we do not have unbiasedness in the small.

Let us see what Dawid’s construction gives when applied to the defensive forecasting strategy constructed from the mixture of $S_n^{\epsilon, I}(p)$ and $S_n^{-\epsilon, I}(p)$, as described

above, over different ϵ and different I ; we will assume not only that the test functions I cover all $[0, 1]$ but also that each point $p \in [0, 1]$ is covered by arbitrarily narrow (concentrated in a small neighborhood of p) test functions. It is clear that we will inevitably have $p_n \rightarrow 0.5$ if p_n are produced by the defensive forecasting strategy and y_n are produced by Dawid’s construction. On the other hand, since all test functions I are continuous and so cannot sharply distinguish between the cases $p_n < 0.5$ and $p_n \geq 0.5$, we do not have any contradiction: neither the test functions nor any observer who can only measure the p_n with a finite precision can detect the lack of unbiasedness in the small.

In this paper we are only interested in unbiasedness in the small when the test functions I are required to be continuous. Dawid’s construction shows that unbiasedness in the small is impossible to achieve if I are allowed to be indicator functions of intervals (such as $[0, 0.5]$ and $[0.5, 1]$). To achieve unbiasedness in the small in this stronger sense, randomization appears necessary (see, e.g., [18]). It is interesting that already a little bit of randomization suffices, as explained in [5].

5 SIMPLIFIED ALGORITHM

Let us assume first that objects are absent, $|\mathbf{X}| = 1$. It was observed empirically that the performance of defensive forecasting strategies with a fixed ϵ does not depend on ϵ much (provided it is not too large; e.g., in the above calculations we assumed $\epsilon \leq 0.5$). This suggests letting $\epsilon \rightarrow 0$ (in particular, we will assume that $\epsilon \ll n^{-2}$). As the test functions I we will take Gaussian bells I_j with standard deviation $\sigma > 0$ located densely and uniformly in the interval $[0, 1]$. Letting \approx stand for approximate equality and using the shorthand $\sum_{\pm} f(\pm) := f(+) + f(-)$, we obtain:

$$\begin{aligned} S_n(p) &= \sum_{\pm} \sum_j (\pm\epsilon) I_j(p) \prod_{i=1}^{n-1} (1 \pm \epsilon I_j(p_i)(y_i - p_i)) \\ &= \sum_{\pm} \sum_j (\pm\epsilon) I_j(p) \exp \left(\sum_{i=1}^{n-1} \ln(1 \pm \epsilon I_j(p_i)(y_i - p_i)) \right) \\ &\approx \sum_{\pm} \sum_j (\pm\epsilon) I_j(p) \exp \left(\pm\epsilon \sum_{i=1}^{n-1} I_j(p_i)(y_i - p_i) \right) \\ &\approx \sum_{\pm} \sum_j (\pm\epsilon) I_j(p) \left(1 \pm \epsilon \sum_{i=1}^{n-1} I_j(p_i)(y_i - p_i) \right) \\ &= \sum_{\pm} \sum_j (\pm\epsilon) I_j(p) \left(\pm\epsilon \sum_{i=1}^{n-1} I_j(p_i)(y_i - p_i) \right) \end{aligned}$$

$$\begin{aligned}
&\propto \sum_j I_j(p) \sum_{i=1}^{n-1} I_j(p_i)(y_i - p_i) \\
&= \sum_{i=1}^{n-1} K(p, p_i)(y_i - p_i), \tag{8}
\end{aligned}$$

where $K(p, p_i)$ is the Mercer kernel

$$K(p, p_i) := \sum_j I_j(p)I_j(p_i).$$

This Mercer kernel can be approximated by

$$\begin{aligned}
&\int_0^1 \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(t-p)^2}{2\sigma^2}\right) \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(t-p_i)^2}{2\sigma^2}\right) dt \\
&\propto \int_0^1 \exp\left(-\frac{(t-p)^2 + (t-p_i)^2}{2\sigma^2}\right) dt \\
&\approx \int_{-\infty}^{\infty} \exp\left(-\frac{(t-p)^2 + (t-p_i)^2}{2\sigma^2}\right) dt.
\end{aligned}$$

As a function of p , the last expression is proportional to the density of the sum of two Gaussian random variables of variance σ^2 ; therefore, it is proportional to

$$\exp\left(-\frac{(p-p_i)^2}{4\sigma^2}\right).$$

To get an idea of the properties of this forecasting strategy, which we call the *K29* strategy (or algorithm), we run it and the Laplace forecasting strategy ($p_n := (k+1)/(n+1)$, where k is the number of 1s observed so far) on a randomly generated bit sequence of length 1000 (with the probability of 1 equal to 0.5). A zero point p_n of S_n was found using the simple bisection procedure (see, e.g., [9], §§9.2–9.4, for more sophisticated methods): (a) start with the interval $[0, 1]$; (b) let p be the mid-point of the current interval; (c) if $S_n(p) > 0$, remove the left half of the current interval; otherwise, remove its right half; (d) go to (b). We did 10 iterations, after which the mid-point of the remaining interval was output as p_n . Notice that the values $S_n(0)$ and $S_n(1)$ did not have to be tested. Our program was written in MATLAB, Version 7, and the initial state of the random number generator was set to 0.

Figure 1 shows that the probabilities output by the K29 ($\sigma = 0.01$) and Laplace forecasting strategies are almost indistinguishable. To see that these two forecasting strategies can behave very differently, we complemented the 1000 bits generated as described above with 1000 0s followed by 1000 1s. The result is shown in Figure 2. The K29 strategy detects that the probability p of 1 changes after the 1000th round, and fairly quickly moves down. When the probability changes

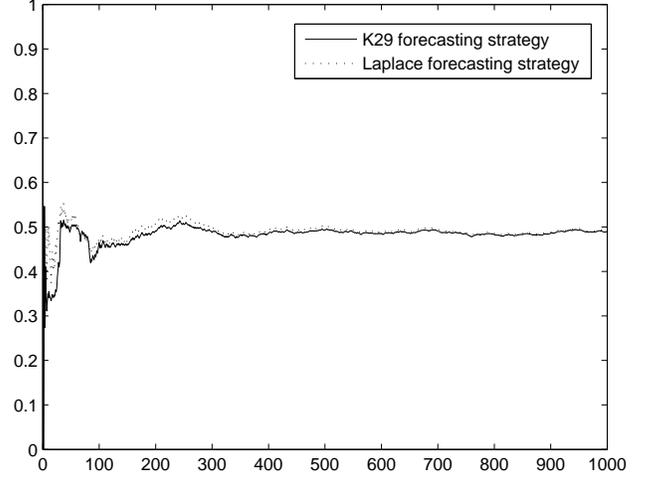


Figure 1: The First 1000 Probabilities Output by the K29 ($\sigma = 0.01$) and Laplace Forecasting Strategies on a Randomly Generated Bit Sequence

again after the 2000th round, K29 starts moving toward $p = 1$, but interestingly, hesitates around the line $p = 0.5$, as if expecting the process to reverse to the original probability of 1.

The Mercer kernel

$$K(p, p_i) = \exp\left(-\frac{(p-p_i)^2}{4\sigma^2}\right)$$

used in these experiments is known in machine learning as the Gaussian kernel (in the usual parameterization $4\sigma^2$ is replaced by $2\sigma^2$ or c); however, many other Mercer kernels also give reasonable results.

If we start from test functions I depending on the object, instead of (8) we will arrive at the expression

$$S_n(p) = \sum_{i=1}^{n-1} K((p, x_n), (p_i, x_i))(y_i - p_i), \tag{9}$$

where K is a Mercer kernel on the squared product $([0, 1] \times \mathbf{X})^2$. There are standard ways of constructing such Mercer kernels from Mercer kernels on $[0, 1]^2$ and \mathbf{X}^2 (see, e.g., the description of tensor products and direct sums in [13, 11]). For S_n to be continuous, we have to require that K be *forecast-continuous* in the following sense: for all $x \in \mathbf{X}$ and all $(p', x') \in [0, 1] \times \mathbf{X}$, $K((p, x), (p', x'))$ is continuous as a function of p . The overall procedure can be summarized as follows.

K29 ALGORITHM

Parameter: forecast-continuous Mercer kernel K on $([0, 1] \times \mathbf{X})^2$

FOR $n = 1, 2, \dots$
 Read $x_n \in \mathbf{X}$.

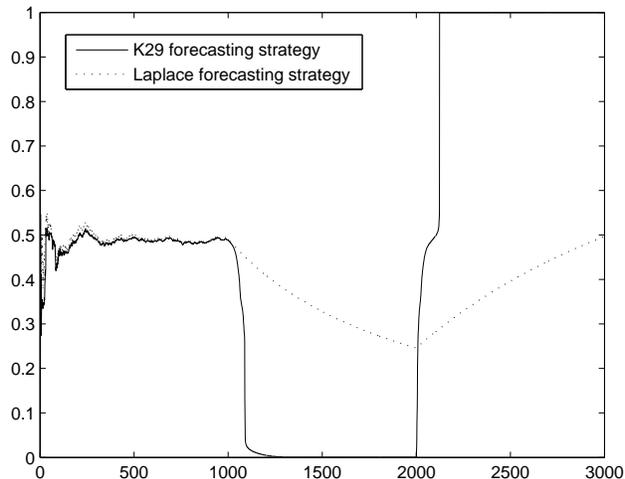


Figure 2: The Probabilities Output by the K29 ($\sigma = 0.01$) and Laplace Forecasting Strategies on a Randomly Generated Sequence of 1000 Bits Followed by 1000 0s and 1000 1s

Define $S_n(p)$ as per (9).
Output any root p of $S_n(p) = 0$ as p_n ;
if there are no roots, $p_n := (1 + \text{sign}(S_n))/2$.
Read $y_n \in \{0, 1\}$.

Computer experiments reported in [16] show that the K29 algorithm performs well on a standard benchmark data set. For a theoretical discussion of the K29 algorithm, see [19] (Appendix) and [17].

6 RELATED WORK AND DIRECTIONS OF FURTHER RESEARCH

This paper’s methods connect two areas that have been developing independently so far: probability forecasting and classical probability theory. It appears that, when properly developed, these methods can benefit both areas:

- the powerful machinery of classical probability theory can be used for probability forecasting;
- practical problems of probability forecasting may suggest new laws of probability.

Classical probability theory started from Bernoulli’s weak law of large numbers (1713) and is the subject of countless monographs and textbooks. The original statements of most of its results were for independent random variables, but they were later extended to the martingale framework; the latter was reduced to its game-theoretic core in [12]. The proof of the strong

law of large numbers used in this paper was extracted from Ville’s [14] martingale proof of the law of the iterated logarithm (upper half).

The theory of probability forecasting was a topic of intensive research in meteorology in the 1960s and 1970s; this research is summarized in [2]. Machine learning is still mainly concerned with categorical prediction, but the situation appears to be changing. Probability forecasting using Bayesian networks is a mature field; the literature devoted to probability forecasting using decision trees and to calibrating other algorithms is also fairly rich. So far, however, the field of probability forecasting has been developing without any explicit connections with classical probability theory.

Defensive forecasting is indirectly related, in a sense dual, to prediction with expert advice (reviewed in [15], §4) and its special case, Bayesian prediction. In prediction with expert advice one starts with a given loss function and tries to make predictions that lead to a small loss as measured by that loss function. In defensive forecasting, one starts with a law of probability and then makes predictions such that this law of probability is satisfied. So the choice of the law of probability when designing the forecasting strategy plays a role analogous to the choice of the loss function in prediction with expert advice.

In prediction with expert advice one combines a pool of potentially promising forecasting strategies to obtain a forecasting strategy that performs not much worse than the best strategies in the pool. In defensive forecasting one combines strategies for Skeptic (such as the strategies corresponding to different test functions I and different $\pm\epsilon$ in §4) to obtain one strategy achieving an interesting goal (such as unbiasedness in the small); a strategy for Forecaster is then obtained using Theorem 1. The possibility of mixing strategies for Skeptic is as fundamental in defensive forecasting as the possibility of mixing strategies for Forecaster in prediction with expert advice.

This paper continues the work started by Foster and Vohra [4] and later developed in, e.g., [6, 10, 18] (the last paper replaces the von Mises–style framework of the previous papers with a martingale framework, as in this paper). The approach of this paper is similar to that of the recent paper [5], which also considers deterministic forecasting strategies and continuous test functions for unbiasedness in the small.

The main difference of this paper’s approach from the bulk of work in learning theory is that we do not make any assumptions about Reality’s strategy.

The following directions of further research appear to us most important:

- extending Theorem 1 to other forecasting protocols (such as multi-label classification) and designing efficient algorithms for finding the corresponding p_n ;
- exploring forecasting strategies corresponding to: (a) Hoeffding’s inequality, (b) the central limit theorem, (c) the law of the iterated logarithm (all we did in this paper was to slightly extend the strong law of large numbers and then use it for probability forecasting).

Acknowledgments

We are grateful to the participants of the PAS-CAL workshop “Notions of complexity: information-theoretic, computational and statistical approaches” (October 2004, EURANDOM) who commented on this work and to the anonymous referees for useful suggestions. This work was partially supported by BBSRC (grant 111/BIO14428), EPSRC (grant GR/R46670/01), MRC (grant S505/65), Royal Society, and, especially, the Superrobust Computation Project (Graduate School of Information Science and Technology, University of Tokyo).

References

- [1] A. Philip Dawid. Self-calibrating priors do not exist: Comment. *Journal of the American Statistical Association*, 80:340–341, 1985. This is a contribution to the discussion in [8].
- [2] A. Philip Dawid. Probability forecasting. In Samuel Kotz, Norman L. Johnson, and Campbell B. Read, editors, *Encyclopedia of Statistical Sciences*, volume 7, pages 210–218. Wiley, New York, 1986.
- [3] Joseph L. Doob. *Stochastic Processes*. Wiley, New York, 1953.
- [4] Dean P. Foster and Rakesh V. Vohra. Asymptotic calibration. *Biometrika*, 85:379–390, 1998.
- [5] Sham M. Kakade and Dean P. Foster. Deterministic calibration and Nash equilibrium. In John Shawe-Taylor and Yoram Singer, editors, *Proceedings of the Seventeenth Annual Conference on Learning Theory*, volume 3120 of *Lecture Notes in Computer Science*, pages 33–48, Heidelberg, 2004. Springer.
- [6] Ehud Lehrer. Any inspection is manipulable. *Econometrica*, 69:1333–1347, 2001.
- [7] Allan H. Murphy and Edward S. Epstein. Verification of probabilistic predictions: a brief review. *Journal of Applied Meteorology*, 6:748–755, 1967.
- [8] David Oakes. Self-calibrating priors do not exist (with discussion). *Journal of the American Statistical Association*, 80:339–342, 1985.
- [9] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C*. Cambridge University Press, Cambridge, second edition, 1992.
- [10] Alvaro Sandroni, Rann Smorodinsky, and Rakesh V. Vohra. Calibration with many checking rules. *Mathematics of Operations Research*, 28:141–153, 2003.
- [11] Bernhard Schölkopf and Alexander J. Smola. *Learning with Kernels*. MIT Press, Cambridge, MA, 2002.
- [12] Glenn Shafer and Vladimir Vovk. *Probability and Finance: It’s Only a Game!* Wiley, New York, 2001.
- [13] Vladimir N. Vapnik. *Statistical Learning Theory*. Wiley, New York, 1998.
- [14] Jean Ville. *Etude critique de la notion de collectif*. Gauthier-Villars, Paris, 1939.
- [15] Vladimir Vovk. Competitive on-line statistics. *International Statistical Review*, 69:213–248, 2001.
- [16] Vladimir Vovk. Defensive forecasting for a benchmark data set, The Game-Theoretic Probability and Finance project, <http://probabilityandfinance.com>, Working Paper #9, September 2004.
- [17] Vladimir Vovk. Non-asymptotic calibration and resolution, The Game-Theoretic Probability and Finance project, <http://probabilityandfinance.com>, Working Paper #11, November 2004.
- [18] Vladimir Vovk and Glenn Shafer. Good randomized sequential probability forecasting is always possible, The Game-Theoretic Probability and Finance project, <http://probabilityandfinance.com>, Working Paper #7, June 2003 (revised September 2004).
- [19] Vladimir Vovk, Akimichi Takemura, and Glenn Shafer. Defensive forecasting, The Game-Theoretic Probability and Finance project, <http://probabilityandfinance.com>, Working Paper #8, September 2004. This is a fuller version of the current paper, with an appendix added.