


Defensive forecasting for linear protocols

Vladimir Vovk, Ilya Nouretdinov,
Akimichi Takemura, and Glenn Shafer



The Game-Theoretic Probability and Finance Project

Working Paper #10

First posted February 3, 2005. Last revised September 24, 2005.

Project web site:
<http://www.probabilityandfinance.com>

Abstract

We consider a general class of forecasting protocols, called “linear protocols”, and discuss several important special cases, including multi-class forecasting. Forecasting is formalized as a game between three players: Reality, whose role is to generate observations; Forecaster, whose goal is to predict the observations; and Skeptic, who tries to make money on any lack of agreement between Forecaster’s predictions and the actual observations. Our main mathematical result is that for any continuous strategy for Skeptic in a linear protocol there exists a strategy for Forecaster that does not allow Skeptic’s capital to grow. This result is a meta-theorem that allows one to transform any continuous law of probability in a linear protocol into a forecasting strategy whose predictions are guaranteed to satisfy this law. We apply this meta-theorem to a weak law of large numbers in Hilbert spaces to obtain a version of the K29 prediction algorithm for linear protocols and show that this version also satisfies the attractive properties of proper calibration and resolution under a suitable choice of its kernel parameter, with no assumptions about the way the data is generated.

Contents

1	Introduction	1
2	Forecasting as a game	1
3	Linear protocol	4
4	Meta-theorem	5
5	A weak law of large numbers in Hilbert space	6
6	The K29 strategy and its properties	8
7	Further research	12
	References	12
A	Zeros of vector fields	14
B	Tensor product	15

1 Introduction

In [14] we suggested a new methodology for designing forecasting strategies. Considering only the simplest case of binary forecasting, we showed that any constructive, in the sense explained below, law of probability can be translated into a forecasting strategy that satisfies this law. In this paper this result is extended to a general class of protocols including multi-class forecasting. In proposing this approach to forecasting we were inspired by [4] and papers further developing [4], although our methods and formal results appear to be completely different.

Whereas the meta-theorem stated in [14] is mathematically trivial, the generalization considered in this paper is less so, depending on the Schauder-Tikhonov fixed-point theorem. Our general meta-theorem is stated in §4 and proved in §4 and Appendix A. The general forecasting protocols covered by this result are introduced and discussed in §§2–3.

In [14] we demonstrated the value of the meta-theorem by applying it to the strong law of large numbers, obtaining from it a kernel forecasting strategy which we called K29. The derivation, however, was informal, involving heuristic transitions to a limit, and this made it impossible to state formally any properties of K29. In this paper we deduce K29 in a much more direct way from the weak law of large numbers and state its properties. (For binary forecasting, this was also done in [13], and the reader might prefer to read that paper first.) The weak law of large numbers is stated and proved in §5, and K29 is derived and studied in §6.

We call the approach to forecasting using our meta-theorem “defensive forecasting”: Forecaster is trying to defend himself when playing against Skeptic. The justification of this approach given in this paper and in [13] is K29’s properties of proper calibration and resolution. Another justification, in a sense the ultimate justification of any forecasts, is given in [12]: defensive forecasts lead to good decisions; this result, however, is obtained in [12] for rather simple decision problems requiring only binary forecasts, and its extensions will require this paper’s results or their generalizations.

The exposition of probability theory needed for this paper is given in [9]. The standard exposition is based on Kolmogorov’s measure-theoretic axioms of probability, whereas [9] states several key laws of probability in terms of a game between the forecaster, the reality, and a third player, the skeptic. The game-theoretic laws of probability in [9] are constructive in that we explicitly construct computable winning strategies for the forecaster in various games of forecasting.

2 Forecasting as a game

Following [9] and [14] we consider the following general forecasting protocol:

FORECASTING GAME 1

Players: Reality, Forecaster, Skeptic

Parameters: \mathbf{X} (*data space*), \mathbf{Y} (*observation space*), \mathbf{F} (*Forecaster's move space*), \mathbf{S} (*Skeptic's move space*), $\lambda : \mathbf{S} \times \mathbf{F} \times \mathbf{Y} \rightarrow \mathbb{R}$ (*Skeptic's gain function and Forecaster's loss function*)

Protocol:

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

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

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

Forecaster announces $f_n \in \mathbf{F}.$

Skeptic announces $s_n \in \mathbf{S}.$

Reality announces $y_n \in \mathbf{Y}.$

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

END FOR

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 ($\mathcal{K}_n \geq 0$). We will say that x_n are the *data*, y_n are the *observations*, and f_n are the *forecasts*. In applications, the datum x_n will contain all available information deemed useful in forecasting y_n .

Book [9] contains several results (game-theoretic versions of limit theorems of probability theory) of the following form: Skeptic has a strategy that guarantees that either a property of agreement between the forecasts f_n and observations y_n is satisfied or Skeptic becomes very rich (without risking bankruptcy, according to the protocol). All specific strategies considered in [9] have computable versions. According to Brouwer's principle (see, e.g., §1 of [10] for a recent review of the relevant literature) they must be automatically continuous; in any case, their continuity can be checked directly. In [14] we showed that, under a special choice of the players' move spaces and Skeptic's gain function λ , for any continuous strategy for Skeptic Forecaster has a strategy that guarantees that Skeptic's capital never increases when he plays that strategy. Therefore, Forecaster has strategies that ensure various properties of agreement between the forecasts and the observations.

The purpose of this paper is to extend the result of [14] to a wide class of Skeptic's gain functions λ . But first we consider several important special cases of Forecasting Game 1.

Binary forecasting

The simplest non-trivial case, considered in [14], is where $\mathbf{Y} = \{0, 1\}$, $\mathbf{F} = [0, 1]$, $\mathbf{S} = \mathbb{R}$, and

$$\lambda(s_n, f_n, y_n) = s_n(y_n - f_n). \quad (1)$$

Intuitively, Forecaster gives probability forecasts for y_n : f_n is his subjective probability that $y_n = 1$. The operational interpretation of f_n is that it is the

price that Forecaster charges for a ticket that will pay y_n at the end of the n th round of the game; s_n is the number (positive, zero, or negative) of such tickets that Skeptic chooses to buy.

Bounded regression

This is the most straightforward extension of binary forecasting, considered in [9], §3.2. The move spaces are $\mathbf{Y} = \mathbf{F} = [A, B]$, where A and B are two constants, and $\mathbf{S} = \mathbb{R}$; the gain function is, as before, (1). This protocol allows one to prove a strong law of large numbers ([9], Proposition 3.3) and a simple one-sided law of the iterated logarithm ([9], Corollary 5.1).

Multi-class forecasting

Another extension of binary forecasting is the protocol where \mathbf{Y} is a finite set, \mathbf{F} is the set of all probability distributions on \mathbf{Y} , \mathbf{S} is the set of all real-valued functions on \mathbf{Y} , and

$$\lambda(s_n, f_n, y_n) = s_n(y_n) - \int s_n \, df_n.$$

The intuition behind Skeptic's move s_n is that Skeptic buys the ticket which pays $s_n(y_n)$ after y_n is announced; he is charged $\int s_n \, df_n$ for this ticket.

The binary forecasting protocol is "isomorphic" to the special case of this protocol where $\mathbf{Y} = \{0, 1\}$: Forecaster's move f_n in the binary forecasting protocol is represented by the probability distribution f'_n on $\{0, 1\}$ assigning weight f_n to $\{1\}$ and Skeptic's move s_n in the binary forecasting protocol is represented by any function s'_n on $\{0, 1\}$ such that $s'_n(1) - s'_n(0) = s_n$. The isomorphism between these two protocols follows from

$$\begin{aligned} s'_n(y_n) - \int s'_n \, df'_n &= s'_n(y_n) - s'_n(1)f_n - s'_n(0)(1 - f_n) \\ &= s'_n(y_n) - s'_n(0) - s_n f_n = s_n(y_n - f_n) \end{aligned}$$

(remember that $y_n \in \{0, 1\}$).

Bounded mean-variance forecasting

In this protocol, $\mathbf{Y} = [A, B]$, where A and B are again two constants, $\mathbf{F} = \mathbf{S} = \mathbb{R}^2$, and

$$\lambda(s_n, f_n, y_n) = \lambda((M_n, V_n), (m_n, v_n), y_n) = M_n(y_n - m_n) + V_n((y_n - m_n)^2 - v_n).$$

Intuitively, Forecaster is asked to forecast y_n with a number m_n and also forecast the accuracy $(y_n - m_n)^2$ of his first forecast with a number v_n . This protocol, although usually without the restriction $y_n \in [A, B]$, is used extensively in [9] (e.g., in Chaps. 4 and 5).

An equivalent representation of this protocol is $\mathbf{Y} = \{(t, t^2) \mid t \in [A, B]\}$, $\mathbf{F} = \mathbf{S} = \mathbb{R}^2$ and

$$\lambda(s_n, f_n, y_n) = \lambda((s'_n, s''_n), (f'_n, f''_n), (t_n, t_n^2)) = s'_n(t_n - f'_n) + s''_n(t_n^2 - f''_n).$$

The equivalence of the two representations can be seen as follows: Reality's move (x_n, t_n) in the first representation corresponds to $(x_n, y_n) = (x_n, (t_n, t_n^2))$ in the second representation, Forecaster's move (m_n, v_n) in the first representation corresponds to $(f'_n, f''_n) = (m_n, v_n + m_n^2)$ in the second representation, and Skeptic's move (s'_n, s''_n) in the second representation corresponds to $(M_n, V_n) = (s'_n + 2m_n s''_n, s''_n)$ in the first representation. This establishes a bijection between Reality's move spaces, a bijection between Forecaster's move spaces, and a bijection between Skeptic's move spaces in the two representations; Skeptic's gains are also the same in the two representations:

$$\begin{aligned} & s'_n(t_n - f'_n) + s''_n(t_n^2 - f''_n) \\ &= s'_n(t_n - m_n) + s''_n\left(\left((t_n - m_n)^2 + 2(t_n - m_n)m_n + m_n^2\right) - (v_n + m_n^2)\right) \\ &= (s'_n + 2m_n s''_n)(t_n - m_n) + s''_n((t_n - m_n)^2 - v_n). \end{aligned}$$

3 Linear protocol

Forecasting Game 1 is too general to derive results of the kind we are interested in. In this subsection we will introduce a narrower protocol which will still be wide enough to cover all special cases considered so far.

All move spaces are now subsets of a Hilbert space \mathbf{L} (we allow \mathbf{L} to be non-separable or finite-dimensional; in fact, in this paper we emphasize the case where $\mathbf{L} = \mathbb{R}^m$ for some positive integer m). The observation space is a non-empty pre-compact subset $\mathbf{Y} \subset \mathbf{L}$ (we say that a set is *pre-compact* if its closure is compact; if $\mathbf{L} = \mathbb{R}^m$, this is equivalent to it being bounded), Forecaster's move space \mathbf{F} is the whole of \mathbf{L} , and Skeptic's move space \mathbf{S} is also the whole of \mathbf{L} . Skeptic's gain function is

$$\lambda(s_n, f_n, y_n) = \langle s_n, y_n - f_n \rangle_{\mathbf{L}}.$$

Therefore, we consider the following perfect-information game:

FORECASTING GAME 2

Players: Reality, Forecaster, Skeptic

Parameters: \mathbf{X} , \mathbf{L} (Hilbert space), \mathbf{Y} (non-empty pre-compact subset of \mathbf{L})

Protocol:

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

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

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

Forecaster announces $f_n \in \mathbf{L}$.

Skeptic announces $s_n \in \mathbf{L}$.

Reality announces $y_n \in \mathbf{Y}$.

$$\mathcal{K}_n := \mathcal{K}_{n-1} + \langle s_n, y_n - f_n \rangle_{\mathbf{L}}. \quad (2)$$

END FOR

Restriction on Skeptic: Skeptic must choose the s_n so that his capital is always nonnegative no matter how the other players move.

Let us check that the specific protocols considered in the previous section are covered by this *linear protocol* (and for all those protocols \mathbf{L} can be taken finite dimensional, $\mathbf{L} = \mathbb{R}^m$ for some $m \in \{1, 2, \dots\}$). At first sight, even the binary forecasting protocol is not covered, as Forecaster's move space is $\mathbf{F} = [0, 1]$ rather than \mathbb{R} . It is easy to see, however, that Forecaster's move $f_n \notin \overline{\text{co}} \mathbf{Y}$ outside the convex closure $\overline{\text{co}} \mathbf{Y}$ of the observation space (the convex closure $\overline{\text{co}} A$ of a set A is defined to be the intersection of all convex closed sets containing A) is always inadmissible, in the sense that there exists Skeptic's reply s_n making him arbitrarily rich regardless of Reality's move, and so we can as well choose $\mathbf{F} := \overline{\text{co}} \mathbf{Y}$. Indeed, suppose that $f_n \notin \overline{\text{co}} \mathbf{Y}$ in the linear protocol. Since \mathbf{Y} is pre-compact, $\overline{\text{co}} \mathbf{Y}$ is compact ([8], Theorem 3.20(c)). By the Hahn-Banach theorem ([8], Theorem 3.4(b)), there exists a vector $s_n \in \mathbf{L}$ such that

$$\inf_{y \in \mathbf{Y}} \langle s_n, y - f_n \rangle_{\mathbf{L}} > 0.$$

(It would have been sufficient for either $\{f_n\}$ or $\overline{\text{co}} \mathbf{Y}$ to be compact; in fact both are.) Skeptic's move Cs_n can make him as rich as he wishes as C can be arbitrarily large. In what follows, we will usually assume that Forecaster's move space is $\overline{\text{co}} \mathbf{Y}$ and use \mathbf{F} as a shorthand for $\overline{\text{co}} \mathbf{Y}$.

Now it is obvious that the binary forecasting, bounded regression, and bounded mean-variance forecasting (in its second representation) protocols are special cases of the linear protocol (perhaps with $\mathbf{F} = \overline{\text{co}} \mathbf{Y}$). For the multi-class forecasting protocol, we should represent \mathbf{Y} as the vertices

$$y^1 := (1, 0, 0, \dots, 0), \quad y^2 := (0, 1, 0, \dots, 0), \dots, \quad y^m := (0, 0, 0, \dots, 1)$$

of the standard simplex in \mathbb{R}^m , where m is the size of \mathbf{Y} , represent the probability distributions f on \mathbf{Y} as vectors $(f\{y^1\}, \dots, f\{y^m\})$ in \mathbb{R}^m , and represent the real-valued functions s on \mathbf{Y} as vectors $(s(y^1), \dots, s(y^m))$ in \mathbb{R}^m .

4 Meta-theorem

In this section we state the main 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. As in [14], we make Skeptic announce his strategy for each round at the outset of that round rather than announce his strategy for the whole game at the beginning of the game, and we drop all restrictions on Skeptic. Forecaster's move space is restricted to $\mathbf{F} = \overline{\text{co}} \mathbf{Y}$. The resulting perfect-information game is:

FORECASTING GAME 3

Players: Reality, Forecaster, Skeptic

Parameters: \mathbf{X}, \mathbf{L} (Hilbert space), $\mathbf{Y} \subset \mathbf{L}$ (non-empty and pre-compact)

Protocol:

\mathcal{K}_0 is set to a real number.

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

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

Skeptic announces continuous $S_n : \overline{\text{co}} \mathbf{Y} \rightarrow \mathbf{L}$.

Forecaster announces $f_n \in \overline{\text{co}} \mathbf{Y}$.

Reality announces $y_n \in \mathbf{Y}$.

$\mathcal{K}_n := \mathcal{K}_{n-1} + \langle S_n(f_n), y_n - f_n \rangle_{\mathbf{L}}$.

END FOR

Theorem 1 *Forecaster has a strategy in Forecasting Game 3 that ensures $\mathcal{K}_0 \geq \mathcal{K}_1 \geq \mathcal{K}_2 \geq \dots$.*

Proof Fix a round n and Skeptic's move $S_n : \mathbf{F} \rightarrow \mathbf{L}$ (we will refer to S_n as a vector field in \mathbf{F}). Our task is to prove the existence of a point $f_n \in \mathbf{F}$ such that, for all $y \in \mathbf{Y}$, $\langle S_n(f_n), y - f_n \rangle_{\mathbf{L}} \leq 0$.

If for some $f \in \partial \mathbf{F}$ (we use ∂A to denote the boundary of $A \subseteq \mathbf{L}$) the vector $S_n(f)$ is normal and directed exteriorly to \mathbf{F} (in the sense that $\langle S_n(f), y - f \rangle_{\mathbf{L}} \leq 0$ for all $y \in \mathbf{F}$), we can take such f as f_n . Therefore, we assume, without loss of generality, that S_n is never normal and directed exteriorly on $\partial \mathbf{F}$. Then by Lemma 1 in Appendix A there exists f such that $S_n(f) = 0$, and we can take such f as f_n . ■

Remark Notice that Theorem 1 will not become weaker if the first move by Reality (choosing x_n) is removed from each round of the protocol.

5 A weak law of large numbers in Hilbert space

Unfortunately, the usual law of large numbers is not useful for the purpose of designing forecasting strategies (see the discussion in [14]). Therefore, we state a generalized law of large numbers; at the end of this section we will explain connections with the usual law of large numbers. In this section we consider Forecasting Game 2 without the requirement $\mathcal{K}_0 = 1$ and with the restriction on Skeptic dropped. If we fix a strategy for Skeptic and Skeptic's initial capital \mathcal{K}_0 (not necessarily 1 or even a positive number), \mathcal{K}_n defined by (2) becomes a function of Reality's and Forecaster's moves. Such functions will be called *capital processes*.

Let $\Phi : \mathbf{F} \times \mathbf{X} \rightarrow \mathbf{H}$ (as usual, $\mathbf{F} = \overline{\text{co}} \mathbf{Y}$) be a *feature mapping* into a Hilbert space \mathbf{H} ; \mathbf{H} is called the *feature space*. The next theorem uses the notion of tensor product; for details, see Appendix B.

Theorem 2 *The function*

$$\mathcal{K}_n := \left\| \sum_{i=1}^n (y_i - f_i) \otimes \Phi(f_i, x_i) \right\|_{\mathbf{L} \otimes \mathbf{H}}^2 - \sum_{i=1}^n \|y_i - f_i\|_{\mathbf{L}}^2 \|\Phi(f_i, x_i)\|_{\mathbf{H}}^2 \quad (3)$$

is a capital process (not necessarily non-negative) of some strategy for Skeptic.

Proof We start by noticing that

$$\begin{aligned} \mathcal{K}_n - \mathcal{K}_{n-1} &= \left\| \sum_{i=1}^{n-1} (y_i - f_i) \otimes \Phi(f_i, x_i) + (y_n - f_n) \otimes \Phi(f_n, x_n) \right\|_{\mathbf{L} \otimes \mathbf{H}}^2 \\ &\quad - \left\| \sum_{i=1}^{n-1} (y_i - f_i) \otimes \Phi(f_i, x_i) \right\|_{\mathbf{L} \otimes \mathbf{H}}^2 - \|y_n - f_n\|_{\mathbf{L}}^2 \|\Phi(f_n, x_n)\|_{\mathbf{H}}^2 \\ &= 2 \left\langle \sum_{i=1}^{n-1} (y_i - f_i) \otimes \Phi(f_i, x_i), (y_n - f_n) \otimes \Phi(f_n, x_n) \right\rangle_{\mathbf{L} \otimes \mathbf{H}} \\ &= 2 \sum_{i=1}^{n-1} \langle y_i - f_i, y_n - f_n \rangle_{\mathbf{L}} \langle \Phi(f_i, x_i), \Phi(f_n, x_n) \rangle_{\mathbf{H}} \end{aligned}$$

(in the last two equalities we used (18) and (19) from Appendix B). Introducing the notation

$$\mathbf{k}((f, x), (f', x')) := \langle \Phi(f, x), \Phi(f', x') \rangle_{\mathbf{H}}, \quad (4)$$

where $(f, x), (f', x') \in \mathbf{F} \times \mathbf{X}$, we can rewrite the expression for $\mathcal{K}_n - \mathcal{K}_{n-1}$ as

$$\left\langle 2 \sum_{i=1}^{n-1} \mathbf{k}((f_i, x_i), (f_n, x_n)) (y_i - f_i), y_n - f_n \right\rangle_{\mathbf{L}}.$$

Therefore, \mathcal{K}_n is the capital process corresponding to Skeptic's strategy

$$2 \sum_{i=1}^{n-1} \mathbf{k}((f_i, x_i), (f_n, x_n)) (y_i - f_i); \quad (5)$$

this completes the proof. ■

More standard statements of the weak law

In the rest of this section we explain connections of Theorem 2 with more standard statements of the weak law of large numbers; in this part of the paper we will use some notions introduced in [9]. The rest of the paper does not depend on this material, and the reader may wish to skip this subsection.

Let us assume that

$$\mathbf{c}_{\Phi} := \sup_{(f, x) \in \mathbf{F} \times \mathbf{X}} \|\Phi(f, x)\|_{\mathbf{H}} < \infty.$$

We will use the notation $\text{diam}(\mathbf{Y}) := \sup_{y, y' \in \mathbf{Y}} \|y - y'\|_{\mathbf{L}}$; it is clear that $\text{diam}(\mathbf{Y}) < \infty$. For any initial capital \mathcal{K}_0 ,

$$\mathcal{K}_n := \mathcal{K}_0 + \left\| \sum_{i=1}^n (y_i - f_i) \otimes \Phi(f_i, x_i) \right\|_{\mathbf{L} \otimes \mathbf{H}}^2 - \sum_{i=1}^n \|y_i - f_i\|_{\mathbf{L}}^2 \|\Phi(f_i, x_i)\|_{\mathbf{H}}^2$$

is the capital process of some strategy for Skeptic. Suppose a positive integer N (the duration of the game, or the *horizon*) is given in advance and $\mathcal{K}_0 := \text{diam}^2(\mathbf{Y}) \mathbf{c}_{\Phi}^2 N$. Then, in the game lasting N rounds, \mathcal{K}_n is never negative and

$$\mathcal{K}_N \geq \left\| \sum_{i=1}^N (y_i - f_i) \otimes \Phi(f_i, x_i) \right\|_{\mathbf{L} \otimes \mathbf{H}}^2.$$

If we do not believe that Skeptic can increase his capital $1/\delta$ -fold for a small $\delta > 0$ without risking bankruptcy, we should believe that

$$\left\| \sum_{i=1}^N (y_i - f_i) \otimes \Phi(f_i, x_i) \right\|_{\mathbf{L} \otimes \mathbf{H}}^2 \leq \text{diam}^2(\mathbf{Y}) \mathbf{c}_{\Phi}^2 N / \delta,$$

which can be rewritten as

$$\left\| \frac{1}{N} \sum_{i=1}^N (y_i - f_i) \otimes \Phi(f_i, x_i) \right\|_{\mathbf{L} \otimes \mathbf{H}} \leq \text{diam}(\mathbf{Y}) \mathbf{c}_{\Phi} (N\delta)^{-1/2}. \quad (6)$$

In the terminology of [9], *the game-theoretic lower probability of the event (6) is at least $1 - \delta$* .

The game-theoretic version of Bernoulli's law of large numbers is a special case of (6) corresponding to $\Phi(f, x) = 1$, for all f and x , $\mathbf{Y} = \{0, 1\}$, and $|\mathbf{X}| = 1$ (the last two conditions mean that we are considering the binary forecasting protocol without the data); as usual, we assume that f_i are chosen from $\overline{\text{co}} \mathbf{Y} = [0, 1]$. As explained in [9], in combination with the measurability of Skeptic's strategy guaranteeing (6), this implies that the measure-theoretic probability of the event (6) is at least $1 - \delta$, assuming that the y_i are generated by a probability distribution and that each f_i is the conditional probability that $y_i = 1$ given y_1, \dots, y_{i-1} . This measure-theoretic result was proved by Kolmogorov in 1929 (see [5]) and is the origin of the name "K29 strategy".

We will see in the next section that the feature-space version (6) of the weak law of large numbers is much more useful than the standard version for the purpose of forecasting.

6 The K29 strategy and its properties

According to Theorem 1, under the continuity assumption there is a strategy for Forecaster that does not allow \mathcal{K}_n to grow, where \mathcal{K}_n is defined by (3).

Fortunately (but not unusually), this strategy depends on the feature mapping Φ only via the corresponding *kernel* \mathbf{k} defined by (4). The continuity assumption needed is that $\mathbf{k}((f, x), (f', x'))$ should be continuous in f ; such kernels will be called *admissible*. According to (5), the corresponding forecasting strategy, which we will call the *K29 strategy* with parameter \mathbf{k} , is to output, on the n th round, a forecast f_n satisfying

$$S(f_n) := \sum_{i=1}^{n-1} \mathbf{k}((f_i, x_i), (f_n, x_n))(y_i - f_i) = 0$$

(or, if such f_n does not exist, the forecast is chosen to be a point $f_n \in \partial \mathbf{F}$ where $S(f_n)$ is normal and directed exteriorly to \mathbf{F}).

The protocol of this section is essentially that of Forecasting Game 3; as Skeptic ceases to be an active player, it simplifies to:

FOR $n = 1, 2, \dots$:
 Reality announces $x_n \in \mathbf{X}$.
 Forecaster announces $f_n \in \overline{\text{co}} \mathbf{Y}$.
 Reality announces $y_n \in \mathbf{Y}$.
END FOR

Theorem 3 *The K29 strategy guarantees that always*

$$\left\| \sum_{i=1}^n (y_i - f_i) \otimes \Phi(f_i, x_i) \right\|_{\mathbf{L} \otimes \mathbf{H}} \leq \text{diam}(\mathbf{Y}) \mathbf{c}_\Phi \sqrt{n}, \quad (7)$$

where $\mathbf{c}_\Phi := \sup_{(f, x) \in \mathbf{F} \times \mathbf{X}} \|\Phi(f, x)\|_{\mathbf{H}}$ is assumed to be finite.

Proof The K29 strategy ensures that (3) never increases; therefore,

$$\left\| \sum_{i=1}^n (y_i - f_i) \otimes \Phi(f_i, x_i) \right\|_{\mathbf{L} \otimes \mathbf{H}}^2 \leq \sum_{i=1}^n \|y_i - f_i\|_{\mathbf{L}}^2 \|\Phi(f_i, x_i)\|_{\mathbf{H}}^2 \leq \text{diam}^2(\mathbf{Y}) \mathbf{c}_\Phi^2 n. \quad \blacksquare$$

Remark The property (7) is a special case of (6) corresponding to $\delta = 1$; we gave an independent derivation to make our exposition self-contained and to avoid the extra assumptions used in the derivation of (6), such as the horizon being finite and known in advance.

K29 with reproducing kernel Hilbert spaces

A *reproducing kernel Hilbert space* (usually abbreviated to RKHS) is a function space \mathcal{F} on some set Z such that all evaluation functionals $F \in \mathcal{F} \mapsto F(z)$, $z \in Z$, are continuous. We will be interested in RKHS on the Cartesian product $\mathbf{F} \times \mathbf{X}$.

By the Riesz-Fischer theorem, for each $z \in Z$ there exists a function $\mathbf{k}_z \in \mathcal{F}$ such that

$$F(z) = \langle \mathbf{k}_z, F \rangle_{\mathcal{F}}, \quad \forall F \in \mathcal{F}.$$

Let

$$\mathbf{c}_{\mathcal{F}} := \sup_{z \in Z} \|\mathbf{k}_z\|_{\mathcal{F}}; \quad (8)$$

we will be interested in the case $\mathbf{c}_{\mathcal{F}} < \infty$.

The *kernel* of an RKHS \mathcal{F} on Z is

$$\mathbf{k}(z, z') := \langle \mathbf{k}_z, \mathbf{k}_{z'} \rangle_{\mathcal{F}} \quad (9)$$

(equivalently, we could define $\mathbf{k}(z, z')$ as $\mathbf{k}_z(z')$ or as $\mathbf{k}_{z'}(z)$). It is clear that (9) is a special case of the generalization

$$\mathbf{k}(z, z') := \langle \Phi(z), \Phi(z') \rangle_{\mathbf{H}} \quad (10)$$

of (4). In fact, the functions \mathbf{k} that can be represented as (10) are exactly the functions that can be represented as (9); they can be equivalently defined as symmetric positive definite functions on Z^2 (see [13] for a list of references).

A long list of RKHS together with their kernels is given in [2], §7.4. We will only give one example: the Sobolev space \mathcal{S} of absolutely continuous functions F on \mathbb{R} with finite norm

$$\|F\|_{\mathcal{S}} := \sqrt{\int_{-\infty}^{\infty} F^2(z) \, dz + \int_{-\infty}^{\infty} (F'(z))^2 \, dz}; \quad (11)$$

its kernel is

$$\mathbf{k}(z, z') = \frac{1}{2} \exp(-|z - z'|)$$

(see [11] or [2], §7.4, Example 24). From the last equation we can see that $\mathbf{c}_{\mathcal{S}} = 1/\sqrt{2}$.

The following is an easy corollary of Theorem 3.

Theorem 4 *Let \mathcal{F} be an RKHS on $\mathbf{F} \times \mathbf{X}$. The K29 strategy with parameter \mathbf{k} (defined by (9)) ensures*

$$\left\| \sum_{i=1}^n F(f_i, x_i)(y_i - f_i) \right\|_{\mathbf{L}} \leq \text{diam}(\mathbf{Y}) \mathbf{c}_{\mathcal{F}} \|F\|_{\mathcal{F}} \sqrt{n} \quad (12)$$

for each function $F \in \mathcal{F}$, where $\mathbf{c}_{\mathcal{F}}$ is defined by (8).

Proof Let $\Phi : \mathbf{F} \times \mathbf{X} \rightarrow \mathbf{H} := \mathcal{F}$ be defined by $\Phi(z) := \mathbf{k}_z$. Theorem 3 then

implies

$$\begin{aligned}
\left\| \sum_{i=1}^n F(f_i, x_i)(y_i - f_i) \right\|_{\mathbf{L}} &= \left\| \sum_{i=1}^n \langle \mathbf{k}_{f_i, x_i}, F \rangle_{\mathbf{H}} (y_i - f_i) \right\|_{\mathbf{L}} \\
&= \left\| \sum_{i=1}^n ((y_i - f_i) \otimes \mathbf{k}_{f_i, x_i}) F \right\|_{\mathbf{L}} \\
&\leq \left\| \sum_{i=1}^n (y_i - f_i) \otimes \mathbf{k}_{f_i, x_i} \right\|_{\mathbf{L} \otimes \mathcal{F}} \|F\|_{\mathcal{F}} \\
&\leq \text{diam}(\mathbf{Y}) \mathbf{c}_{\mathcal{F}} \|F\|_{\mathcal{F}} \sqrt{n}
\end{aligned}$$

(the second equality follows from Lemma 2 and the first inequality from Lemma 3 in Appendix B). \blacksquare

Calibration and resolution

Two important properties of a forecasting strategy are its calibration and resolution, which we introduce informally. Our discussion in this section extends the discussion in [13], §5, to the case of linear protocols (in particular, to the case of multi-class forecasting). Forecaster’s move space is assumed to be $\mathbf{F} = \overline{\text{co}} \mathbf{Y}$.

We say that the forecasts f_n are *properly calibrated* if, for any $f^* \in \mathbf{F}$,

$$\frac{\sum_{i=1, \dots, n: f_i \approx f^*} y_i}{\sum_{i=1, \dots, n: f_i \approx f^*} 1} \approx f^*$$

provided $\sum_{i=1, \dots, n: f_i \approx f^*} 1$ is not too small. (We shorten $(1/c)v$ to v/c , where v is a vector and $c \neq 0$ is a number.) Proper calibration is only a necessary but far from sufficient condition for good forecasts: for example, a forecaster who ignores the data x_n can be perfectly calibrated, no matter how much useful information x_n contain. (Cf. the discussion in [3].)

We say that the forecasts f_n are *properly calibrated and resolved* if, for any $(f^*, x^*) \in \mathbf{F} \times \mathbf{X}$,

$$\frac{\sum_{i=1, \dots, n: (f_i, x_i) \approx (f^*, x^*)} y_i}{\sum_{i=1, \dots, n: (f_i, x_i) \approx (f^*, x^*)} 1} \approx f^* \tag{13}$$

provided $\sum_{i=1, \dots, n: (f_i, x_i) \approx (f^*, x^*)} 1$ is not too small.

Instead of “crisp” points $(f^*, x^*) \in \mathbf{F} \times \mathbf{X}$ one may consider “fuzzy points” $I : \mathbf{F} \times \mathbf{X} \rightarrow [0, 1]$ such that $I(f^*, x^*) = 1$ and $I(f, x) = 0$ for all (f, x) outside a small neighborhood of (f^*, x^*) . A standard choice would be something like $I := \mathbb{I}_E$, where $E \subseteq \mathbf{F} \times \mathbf{X}$ is a small neighborhood of (f^*, x^*) and \mathbb{I}_E is its indicator function, but we will want I to be continuous (it can, however, be arbitrarily close to \mathbb{I}_E).

Suppose $\mathbf{F} \subseteq \mathbb{R}^m$ and $\mathbf{X} \subseteq \mathbb{R}^l$ for some $m, l \in \{1, 2, \dots\}$. Let (f^*, x^*) be a point in $\mathbf{F} \times \mathbf{X}$; consider a small box $E := \prod_{i=1}^m [a_i, b_i] \times \prod_{j=1}^l [c_j, d_j]$

containing this point, $E \ni (f^*, x^*)$. The indicator $\mathbb{1}_E$ of E can be arbitrarily well approximated by the tensor product

$$I(f_1, \dots, f_m, x_1, \dots, x_l) = \prod_{i=1}^m F_i(f_i) \prod_{j=1}^l G_j(x_j)$$

of some functions F_i and G_j from the Sobolev class (11). Let $\|I\|_{\mathcal{F}}$ be the norm of I in the tensor product \mathcal{F} of $m + l$ copies of \mathcal{S} (see [1], §1.8, for an explicit description of tensor products of RKHS). We can rewrite (12) as

$$\left\| \frac{\sum_{i=1}^n I(f_i, x_i)(y_i - f_i)}{\sum_{i=1}^n I(f_i, x_i)} \right\|_{\mathbf{L}} \leq 2^{-\frac{m+l}{2}} \frac{\text{diam}(\mathbf{Y}) \|I\|_{\mathcal{F}} \sqrt{n}}{\sum_{i=1}^n I(f_i, x_i)} \quad (14)$$

(assuming the denominator $\sum_{i=1}^n I(f_i, x_i)$ is positive); therefore, we can expect proper calibration and resolution in the soft neighborhood I of (f^*, x^*) when

$$\sum_{i=1}^n I(f_i, x_i) \gg \sqrt{n}. \quad (15)$$

7 Further research

The main result of this paper is an existence theorem: we did not show how to compute Forecaster's strategy ensuring $\mathcal{K}_0 \geq \mathcal{K}_1 \geq \dots$. (The latter was easy in the case of binary forecasting considered in [14].) It is important to develop computationally efficient ways to find zeros of vector fields, at least when $\mathbf{L} = \mathbb{R}^m$. There are several popular methods for finding zeros, such as the Newton-Raphson method (see, e.g., [6], Chap. 9), but it would be ideal to have efficient methods that are guaranteed to find a zero (or a near zero) in a prespecified time.

Acknowledgments

This work was partially supported by MRC (grant S505/65), Royal Society, and the Superrobust Computation Project (Graduate School of Information Science and Technology, University of Tokyo). We are grateful to anonymous reviewers for their comments.

References

- [1] Nachman Aronszajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68:337–404, 1950.
- [2] Alain Berlinet and Christine Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Kluwer, Boston, 2004.

- [3] 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.
- [4] Dean P. Foster and Rakesh V. Vohra. Asymptotic calibration. *Biometrika*, 85:379–390, 1998.
- [5] Andrei N. Kolmogorov. Sur la loi des grands nombres. *Atti della Reale Accademia Nazionale dei Lincei. Classe di scienze fisiche, matematiche, e naturali. Rendiconti Serie VI*, 185:917–919, 1929.
- [6] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C*. Cambridge University Press, Cambridge, second edition, 1992.
- [7] Michael Reed and Barry Simon. *Functional Analysis*. Academic Press, New York, 1972.
- [8] Walter Rudin. *Functional Analysis*. McGraw-Hill, Boston, second edition, 1991.
- [9] Glenn Shafer and Vladimir Vovk. *Probability and Finance: It's Only a Game!* Wiley, New York, 2001.
- [10] Viggo Stoltenberg-Hansen and John V. Tucker. Computable and continuous partial homomorphisms on metric partial algebras. *Bulletin of Symbolic Logic*, 9:299–334, 2003.
- [11] Christine Thomas-Agnan. Computing a family of reproducing kernels for statistical applications. *Numerical Algorithms*, 13:21–32, 1996.
- [12] Vladimir Vovk. Defensive forecasting with expert advice, The Game-Theoretic Probability and Finance project, <http://probabilityandfinance.com>, Working Paper #14, June 2005. First posted in May 2005.
- [13] Vladimir Vovk. Non-asymptotic calibration and resolution, The Game-Theoretic Probability and Finance project, <http://probabilityandfinance.com>, Working Paper #13, May 2005. Originally posted as Working Paper #11, November 2004.
- [14] Vladimir Vovk, Akimichi Takemura, and Glenn Shafer. Defensive forecasting, The Game-Theoretic Probability and Finance project, <http://probabilityandfinance.com>, Working Paper #8, September 2004 (revised January 2005). A shorter version is published in the Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics and available electronically at <http://www.gatsby.ucl.ac.uk/aistats/>.

A Zeros of vector fields

The following lemma is the main component of the proof of Theorem 1.

Lemma 1 *Let \mathbf{F} be a compact convex non-empty set in a Hilbert space \mathbf{L} and $S : \mathbf{F} \rightarrow \mathbf{L}$ be a continuous vector field on \mathbf{F} . If at no point of the boundary $\partial\mathbf{F}$ the vector field S is normal and directed exteriorly to \mathbf{F} then there exists $f \in \mathbf{F}$ such that $S(f) = 0$.*

Proof For each $f \in \mathbf{L}$ define $\sigma(f)$ to be the point of \mathbf{F} closest to f . A standard argument (see, e.g., [8], Theorem 12.3) shows that such a point exists: if $d := \inf\{\|y - f\|_{\mathbf{L}} \mid y \in \mathbf{F}\}$, we can take any sequence $y_n \in \mathbf{F}$ with $\|y_n - f\|_{\mathbf{L}} \rightarrow d$ and apply the parallelogram law $\|a - b\|^2 + \|a + b\|^2 = 2\|a\|^2 + 2\|b\|^2$ to obtain

$$\begin{aligned} \|y_m - y_n\|_{\mathbf{L}}^2 &= \|(y_m - f) - (y_n - f)\|_{\mathbf{L}}^2 \\ &= 2\|y_m - f\|_{\mathbf{L}}^2 + 2\|y_n - f\|_{\mathbf{L}}^2 - \|(y_m - f) + (y_n - f)\|_{\mathbf{L}}^2 \\ &= 2\|y_m - f\|_{\mathbf{L}}^2 + 2\|y_n - f\|_{\mathbf{L}}^2 - 4\left\|\frac{y_m + y_n}{2} - f\right\|_{\mathbf{L}}^2 \\ &\leq 2\|y_m - f\|_{\mathbf{L}}^2 + 2\|y_n - f\|_{\mathbf{L}}^2 - 4d^2 \rightarrow 2d^2 + 2d^2 - 4d^2 = 0 \end{aligned}$$

as $m, n \rightarrow \infty$; since \mathbf{L} is complete and \mathbf{F} is closed, $y_n \rightarrow y$ for some $y \in \mathbf{F}$, and it is clear that $\|y - f\|_{\mathbf{L}} = d$. A closest point is indeed unique: if $\|y_1 - f\|_{\mathbf{L}} = \|y_2 - f\|_{\mathbf{L}} = d$ and $y_1 \neq y_2$, the parallelogram law would give

$$\begin{aligned} \left\|\frac{y_1 + y_2}{2} - f\right\|_{\mathbf{L}}^2 &= \frac{1}{4}\|(y_1 - f) + (y_2 - f)\|_{\mathbf{L}}^2 \\ &= \frac{1}{2}\|y_1 - f\|_{\mathbf{L}}^2 + \frac{1}{2}\|y_2 - f\|_{\mathbf{L}}^2 - \frac{1}{4}\|(y_1 - f) - (y_2 - f)\|_{\mathbf{L}}^2 \\ &= d^2 - \frac{1}{4}\|y_1 - y_2\|_{\mathbf{L}}^2 < d^2. \end{aligned} \quad (16)$$

Therefore, the function $\sigma(f)$ is well-defined. It is also continuous: if $\|f - \sigma(f)\|_{\mathbf{L}} = d$ and $f_n \rightarrow f$, then $\|f - \sigma(f_n)\|_{\mathbf{L}} \rightarrow d$ and, analogously to (16),

$$\begin{aligned} d^2 &\leq \left\|\frac{\sigma(f) + \sigma(f_n)}{2} - f\right\|_{\mathbf{L}}^2 = \frac{1}{4}\|(\sigma(f) - f) + (\sigma(f_n) - f)\|_{\mathbf{L}}^2 \\ &= \frac{1}{2}\|\sigma(f) - f\|_{\mathbf{L}}^2 + \frac{1}{2}\|\sigma(f_n) - f\|_{\mathbf{L}}^2 - \frac{1}{4}\|(\sigma(f) - f) - (\sigma(f_n) - f)\|_{\mathbf{L}}^2 \\ &= d^2 + o(1) - \frac{1}{4}\|\sigma(f) - \sigma(f_n)\|_{\mathbf{L}}^2; \end{aligned}$$

therefore, $\sigma(f_n) \rightarrow \sigma(f)$ in \mathbf{L} .

For each $f \in \mathbf{F}$, let $\Sigma(f) := \sigma(f + S(f))$ be the point of \mathbf{F} closest to $f + S(f)$; since both σ and S are continuous, Σ is continuous. By the Schauder-Tikhonov

theorem (see, e.g., [8], Theorem 5.28) there is a point $f \in \mathbf{F}$ such that $\Sigma(f) = f$. If f is an interior point of \mathbf{F} , $\sigma(f + S(f)) = f$ implies $S(f) = 0$, and so the conclusion of the lemma holds. It remains to consider the case $f \in \partial\mathbf{F}$; in fact, we will show that this case is impossible. There exists $y \in \mathbf{F}$ such that $\langle S(f), y - f \rangle_{\mathbf{L}} > 0$ (otherwise, S would have been normal and directed exteriorly to \mathbf{F}), and we find for $t \in (0, 1)$:

$$\begin{aligned} \|(f + S(f)) - ((1-t)f + ty)\|_{\mathbf{L}}^2 &= \|S(f) - t(y - f)\|_{\mathbf{L}}^2 \\ &= \|S(f)\|_{\mathbf{L}}^2 - 2t \langle S(f), y - f \rangle_{\mathbf{L}} + t^2 \|y - f\|_{\mathbf{L}}^2; \end{aligned}$$

for a small enough t this gives

$$\|(f + S(f)) - ((1-t)f + ty)\|_{\mathbf{L}}^2 < \|S(f)\|_{\mathbf{L}}^2,$$

a contradiction. ■

B Tensor product

In this appendix we list several definitions and simple facts about tensor products of Hilbert spaces, in the form used in this paper.

The *tensor product* $\mathbf{L} \otimes \mathbf{H}$ of Hilbert spaces \mathbf{L} and \mathbf{H} is defined in, e.g., [7], §II.4. Briefly, the definition is as follows. The space $\mathbf{L} \otimes \mathbf{H}$ is the subset of the set of bilinear forms $v(l', h')$, $l' \in \mathbf{L}$ and $h' \in \mathbf{H}$, obtained as the completion of the set of all linear combinations of the bilinear forms $l \otimes h$, where $l \in \mathbf{L}$ and $h \in \mathbf{H}$, defined by

$$(l \otimes h)(l', h') := \langle l, l' \rangle_{\mathbf{L}} \langle h, h' \rangle_{\mathbf{H}}; \quad (17)$$

the inner product in $\mathbf{L} \otimes \mathbf{H}$ is determined uniquely by setting

$$\langle l_1 \otimes h_1, l_2 \otimes h_2 \rangle_{\mathbf{L} \otimes \mathbf{H}} := \langle l_1, l_2 \rangle_{\mathbf{L}} \langle h_1, h_2 \rangle_{\mathbf{H}}. \quad (18)$$

In particular, (18) implies

$$\|l \otimes h\|_{\mathbf{L} \otimes \mathbf{H}} = \|l\|_{\mathbf{L}} \|h\|_{\mathbf{H}} \quad (19)$$

for all $l \in \mathbf{L}$ and $h \in \mathbf{H}$.

If $v \in \mathbf{L} \otimes \mathbf{H}$ and $h \in \mathbf{H}$, we define the *product* $vh \in \mathbf{L}$ by the requirement

$$v(l', h) = \langle vh, l' \rangle_{\mathbf{L}}, \quad \forall l' \in \mathbf{L}$$

(the validity of this definition follows from the Riesz-Fischer theorem: all bilinear forms in $\mathbf{L} \otimes \mathbf{H}$ are clearly continuous).

Lemma 2 *For any $l \in \mathbf{L}$ and $h_1, h_2 \in \mathbf{H}$,*

$$(l \otimes h_1)h_2 = \langle h_1, h_2 \rangle_{\mathbf{H}} l. \quad (20)$$

Proof It suffices to prove

$$\langle (l \otimes h_1)h_2, l' \rangle_{\mathbf{L}} = \langle h_1, h_2 \rangle_{\mathbf{H}} \langle l, l' \rangle_{\mathbf{L}},$$

which, by definition, is equivalent to

$$(l \otimes h_1)(l', h_2) = \langle h_1, h_2 \rangle_{\mathbf{H}} \langle l, l' \rangle_{\mathbf{L}}$$

and, therefore, true (cf. (17)). ■

The following lemma is an easy implication of the Cauchy-Schwarz inequality.

Lemma 3 *For any $v \in \mathbf{L} \otimes \mathbf{H}$ and $h \in \mathbf{H}$,*

$$\|vh\|_{\mathbf{L}} \leq \|v\|_{\mathbf{L} \otimes \mathbf{H}} \|h\|_{\mathbf{H}}.$$

Proof We are required to prove, for all $l' \in \mathbf{L}$,

$$\langle vh, l' \rangle_{\mathbf{L}} \leq \|v\|_{\mathbf{L} \otimes \mathbf{H}} \|h\|_{\mathbf{H}} \|l'\|_{\mathbf{L}},$$

i.e.,

$$v(l', h) \leq \|v\|_{\mathbf{L} \otimes \mathbf{H}} \|h\|_{\mathbf{H}} \|l'\|_{\mathbf{L}}.$$

We can assume that $v = l \otimes h'$, for some $l \in \mathbf{L}$ and $h' \in \mathbf{H}$, in which case the last inequality immediately follows from (17), (19), and the Cauchy-Schwarz inequality. ■