


Continuous and randomized defensive forecasting: unified view

Vladimir Vovk



The Game-Theoretic Probability and Finance Project

Working Paper #21

First posted August 19, 2007. Last revised August 23, 2007.

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

Abstract

Defensive forecasting is a method of transforming laws of probability (stated in game-theoretic terms as strategies for Sceptic) into forecasting algorithms. There are two known varieties of defensive forecasting: “continuous”, in which Sceptic’s moves are assumed to depend on the forecasts in a (semi)continuous manner and which produces deterministic forecasts, and “randomized”, in which the dependence of Sceptic’s moves on the forecasts is arbitrary and Forecaster’s moves are allowed to be randomized. This note shows that the randomized variety can be obtained from the continuous variety by smearing Sceptic’s moves to make them continuous.

Contents

1	Introduction	1
2	Continuous defensive forecasting	1
3	Randomized defensive forecasting	3
4	Discussion	8
	References	8

1 Introduction

The continuous variety of defensive forecasting was essentially introduced by Levin [9], but was later rediscovered by Kakade and Foster [7] and Takemura *et al.* [16].

The randomized variety was introduced (in the case of von Mises's version of the game-theoretic approach to probability) by Foster and Vohra [5] and further developed by, among others, Sandroni *et al.* [11]; these papers, however, were only concerned with asymptotic calibration. Non-asymptotic versions of the randomized variety were proposed by Sandroni [10] (based on standard measure-theoretic probability) and Vovk and Shafer [15] (based on game-theoretic probability). Kakade and Foster [7] noticed that some calibration results require very little randomization (this will be an important aspect of our Theorem 2).

This note states two simple results about defensive forecasting, Theorem 1 about the continuous variety and Theorem 2 about the randomized variety. The proof of Theorem 2 is obtained from the proof of Theorem 1 by blurring Sceptic's moves.

In our informal discussions we will be assuming that the set Ω of all possible outcomes is finite, although we will try to make mathematical statements as general as possible. The reader who is only interested in the main ideas might choose to specialize Theorems 1 and 2 and their proofs to the case of finite Ω .

2 Continuous defensive forecasting

Let Ω (the *outcome space*) be a compact (i.e., a compact Hausdorff topological space) equipped with the Baire σ -algebra and $\mathcal{P}(\Omega)$ be the set of all probability measures on Ω equipped with the standard topology (the weak* topology on $\mathcal{P}(\Omega)$ identified with a subset of $C(\Omega)'$ by a Riesz representation theorem, Theorem 7.4.1 in [3]; this is also known as the topology of weak convergence in the case of metrizable Ω). The subset $\mathcal{P}^{\text{fin}}(\Omega)$ of $\mathcal{P}(\Omega)$ consists of all probability measures in $\mathcal{P}(\Omega)$ concentrated on a finite subset of Ω . If Ω is finite, $\mathcal{P}(\Omega) = \mathcal{P}^{\text{fin}}(\Omega)$ can be identified with an $(|\Omega| - 1)$ -dimensional simplex (see below) in Euclidean space equipped with the standard Euclidean distance and topology.

Theorem 1 will be a statement about the following perfect-information game involving three players:

CONTINUOUS GAME

Players: Sceptic, Forecaster, Reality

Protocol:

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

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

Sceptic announces a function $S_n : \Omega \times \mathcal{P}(\Omega) \rightarrow \mathbb{R}$

which is lower semicontinuous in the second argument

and satisfies $\int_{\Omega} S_n(\omega, p)p(d\omega) \leq 0$ for all $p \in \mathcal{P}^{\text{fin}}(\Omega)$.

Forecaster announces $p_n \in \mathcal{P}(\Omega)$.
Reality announces $\omega_n \in \Omega$.
 $\mathcal{K}_n := \mathcal{K}_{n-1} + S_n(\omega_n, p_n)$.

Winner: Forecaster wins if Sceptic's capital \mathcal{K}_n stays bounded.

(For $p \in \mathcal{P}^{\text{fin}}(\Omega)$, the integral $\int_{\Omega} S_n(\omega, p)p(d\omega)$ is interpreted as a sum, and so $S_n(\omega, p)$ is not required to be measurable in ω .)

Intuitively, on each round of the game Forecaster is asked to give a probability forecast p_n for the outcome ω_n to be chosen by Reality. Sceptic is testing the forecasts p_n by gambling against them. Forecaster wins the game if Sceptic does not detect serious disagreement between Forecaster and Reality.

The continuous game is stated here in the form that makes Theorem 1 as strong as possible. In typical applications in prediction with expert advice and algorithmic information theory, Sceptic's move $S_n(\omega, p)$ is lower semi-continuous jointly in $(\omega, p) \in \Omega \times \mathcal{P}(\Omega)$ and measurable in ω ; the condition $\int_{\Omega} S_n(\omega, p)p(d\omega) \leq 0$ is required to hold for all $p \in \mathcal{P}(\Omega)$. Furthermore, there is an important restriction imposed on Sceptic: he must choose S_n so that his capital remains nonnegative ($\mathcal{K}_n \geq 0$) no matter how the other players move (in particular, the function S_n must be bounded below). Theorem 1, however, does not depend on these further assumptions.

The following result was stated (in different terms) by Levin [9].

Theorem 1 *Forecaster has a strategy in the continuous game that guarantees $\mathcal{K}_0 \geq \mathcal{K}_1 \geq \mathcal{K}_2 \geq \dots$.*

In other words, not only Sceptic does not detect serious disagreement between Forecaster and Reality, he does not detect any disagreement at all.

We will reproduce Levin's original proof, as detailed by Gács [6], Section 5; for a different proof (essentially a reference to Ky Fan's minimax theorem, [1], Theorem 11.4) under stronger assumptions, see [14], Section 3.

A set v_1, \dots, v_M of points in a Euclidean (or topological vector) space is *affinely independent* if, for all real numbers $\lambda_1, \dots, \lambda_M$,

$$\sum_{m=1}^M \lambda_m v_m = 0 \text{ and } \sum_{m=1}^M \lambda_m = 0 \text{ imply } \lambda_1 = \dots = \lambda_M = 0.$$

The convex hull of such v_1, \dots, v_M , denoted $\text{co}(v_1, \dots, v_M)$, is called a *simplex* or, more fully, an $(M - 1)$ -*dimensional simplex*. The proof of Theorem 1 will use the following result due to Knaster, Kuratowski, and Mazurkiewicz ([8]; see also [1], Theorem 11.2).

KKM Theorem *Let F_1, \dots, F_M be closed subsets of a simplex $\text{co}(v_1, \dots, v_M)$. Suppose that for all $1 \leq k \leq M$ and $1 \leq m_1 \leq \dots \leq m_k \leq M$ we have*

$$\text{co}(v_{m_1}, \dots, v_{m_k}) \subseteq F_{m_1} \cup \dots \cup F_{m_k}.$$

Then $F_1 \cap \dots \cap F_M \neq \emptyset$.

Proof of Theorem 1: Fix a round n of the game and set $S := S_n$. For every $\omega \in \Omega$, let F_ω be the closed set

$$F_\omega := \{p \in \mathcal{P}(\Omega) \mid S(\omega, p) \leq 0\}.$$

It suffices to show that for every finite set of points $\omega_1, \dots, \omega_M$ we have

$$F_{\omega_1} \cap \dots \cap F_{\omega_M} \neq \emptyset. \quad (1)$$

Indeed, the compactness of Ω implies the compactness of $\mathcal{P}(\Omega)$ (combine Alaoglu's theorem, Problem 9 in Section 6.1 of [3], with the weak* closeness of $\mathcal{P}(\Omega)$ in $C(\Omega)'$, following from [3], Theorems 7.1.5 and 2.6.3). Therefore, if every finite subset of the family $\{F_\omega \mid \omega \in \Omega\}$ of closed sets has a non-empty intersection, then the whole family has a nonempty intersection, and any of the measures in this intersection can be taken as p_n .

To show (1), let $\mathcal{P}(\omega_1, \dots, \omega_M)$ be the set of probability measures concentrated on $\{\omega_1, \dots, \omega_M\}$. If $p \in \mathcal{P}(\omega_1, \dots, \omega_M)$, the inequality $\int S(\omega, p)p(d\omega) \leq 0$ implies $S(\omega_m, p) \leq 0$ for some $m \in \{1, \dots, M\}$. Hence $\mathcal{P}(\omega_1, \dots, \omega_M) \subseteq F_{\omega_1} \cup \dots \cup F_{\omega_M}$, and the same holds for every subset of the indices $\{1, \dots, M\}$. The KKM theorem now implies (1). ■

3 Randomized defensive forecasting

Let $\mathcal{P}^{\text{fin}}(\mathcal{P}(\Omega))$ be the set of all probability measures on $\mathcal{P}(\Omega)$ concentrated on a finite subset of $\mathcal{P}(\Omega)$. For each $P \in \mathcal{P}^{\text{fin}}(\mathcal{P}(\Omega))$, let $D(P) \subseteq \mathcal{P}(\Omega)$ be the smallest finite set in $\mathcal{P}(\Omega)$ of P -probability one.

Our result about randomized defensive forecasting concerns the following perfect-information game involving four players:

RANDOMIZED GAME

Players: Sceptic, Forecaster, Reality, Random Number Generator

Protocol:

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

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

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

Sceptic announces a function $S_n : \Omega \times \mathcal{P}(\Omega) \rightarrow \mathbb{R}$
which is continuous in the first argument $\omega \in \Omega$
and satisfies $\int_\Omega S_n(\omega, p)p(d\omega) \leq 0$ for all $p \in \mathcal{P}(\Omega)$.

Forecaster announces $P_n \in \mathcal{P}^{\text{fin}}(\mathcal{P}(\Omega))$.

Reality announces $\omega_n \in \Omega$.

Forecaster announces a function $f_n : \mathcal{P}(\Omega) \rightarrow \mathbb{R}$ such that $\int_{\mathcal{P}(\Omega)} f_n dP_n \leq 0$.

Random Number Generator announces $p_n \in D(P_n)$.

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

$\mathcal{F}_n := \mathcal{F}_{n-1} + f_n(p_n)$.

Restriction on Sceptic: Sceptic must choose S_n (continuous, and so Baire measurable, in its first argument) so that his capital remains nonnegative

($\mathcal{K}_n \geq 0$) no matter how the other players move (in particular, the function S_n must be bounded below).

Restriction on Forecaster: Forecaster must choose his moves so that his capital remains nonnegative ($\mathcal{F}_n \geq 0$) no matter how the other players move.

Winner: Forecaster wins if either (i) his capital \mathcal{F}_n tends to infinity or (ii) Sceptic's capital \mathcal{K}_n stays bounded.

(Since $\int_{\mathcal{P}(\Omega)} f_n dP_n = \int_{\mathcal{D}(P_n)} f_n dP_n$ is a sum, its existence does not depend on the measurability of f_n . However, by the Tietze–Urysohn theorem, Theorem 2.1.8 in [4], f_n can be chosen continuous and, therefore, Baire measurable; the Tietze–Urysohn theorem is applicable since every compact is normal, [4], Theorem 3.1.9.)

Forecaster is now allowed to randomize, and it is Random Number Generator who picks the actual forecast p_n from Forecaster's randomized forecast P_n . As before, Sceptic is testing the forecasts p_n by gambling against them. To make sure that Random Number Generator performs his duty of producing random-looking p_n , Forecaster is allowed to gamble against Random Number Generator's choices. Forecaster wins the game if he either discredits Random Number Generator or Sceptic does not detect serious disagreement between the forecasts and the outcomes.

In the case of finite Ω , the only restriction on Sceptic's move S_n is $\int_{\Omega} S_n(\omega, p) p(d\omega) \leq 0$, $\forall p \in \mathcal{P}(\Omega)$. We will see that in this case Theorem 2 will remain true even if f_n is required to be a linear function on the simplex $\mathcal{P}(\Omega)$.

The following is the randomized counterpart of Theorem 1.

Theorem 2 *For any $\epsilon > 0$ and any sequence $\mathcal{A}_1, \mathcal{A}_2, \dots$ of open covers of the outcome space Ω , Forecaster has a strategy in the randomized game that guarantees:*

- $\mathcal{K}_n \leq (1 + \epsilon)\mathcal{F}_n$ for each n ;
- $\mathcal{D}(P_n)$ lies completely in one element of \mathcal{A}_n ;
- $|\mathcal{D}(P_n)| \leq |\Omega|$.

The last item, $|\mathcal{D}(P_n)| \leq |\Omega|$, is of interest only in the case of finite Ω : it holds trivially when Ω is infinite.

Before discussing the intuition behind Theorem 2 we restate the second item in a more intuitive form assuming that Ω is finite and dist is the Euclidean distance on the simplex $\mathcal{P}(\Omega)$. (More generally, Ω can be assumed a compact metric space and dist be, e.g., the Prokhorov metric on $\mathcal{P}(\Omega)$; see, e.g., [2], Appendix III, Theorem 6.)

Corollary 1 *Suppose Ω is finite (or a metric compact). For any $\epsilon > 0$ and any sequence $\epsilon_1, \epsilon_2, \dots$ of positive real numbers, Forecaster has a strategy in the randomized game that guarantees:*

- $\mathcal{K}_n \leq (1 + \epsilon)\mathcal{F}_n$ for each n ;
- the diameter of $\mathsf{D}(P_n)$ is at most ϵ_n :

$$\text{diam } \mathsf{D}(P_n) := \sup_{p,q \in \mathsf{D}(P_n)} \text{dist}(p,q) = \max_{p,q \in \mathsf{D}(P_n)} \text{dist}(p,q) \leq \epsilon_n;$$

- $|\mathsf{D}(P_n)| \leq |\Omega|$.

The condition $\mathcal{K}_n \leq (1 + \epsilon)\mathcal{F}_n$ says that Forecaster can guarantee $\mathcal{F}_n \geq \mathcal{K}_n$ to any approximation required, i.e., every pound gained by Sceptic can be attributed to the poor performance of Random Number Generator. The condition $\text{diam } \mathsf{D}(P_n) \leq \epsilon_n$ shows that already a tiny amount of randomization is sufficient; as already mentioned, a similar observation was made by Kakade and Foster [7].

Proof of Theorem 2: We will repeatedly use the fact that $\mathcal{P}(\Omega)$ is paracompact ([4], Theorem 5.1.1). The stronger condition that $\mathcal{P}(\Omega)$ is compact will only be used in a reference to Theorem 1.

Fix a round n of the game. Let $\delta > 0$ be a small constant (how small will be determined later). For each $p \in \mathcal{P}(\Omega)$ set

$$A_p := \left\{ q \in \mathcal{P}(\Omega) \mid \int_{\Omega} S_n(\omega, p)q(d\omega) < \delta \right\}; \quad (2)$$

notice that $p \in A_p$ and that A_p is an open set. Let \mathcal{B} be any open star refinement of \mathcal{A}_n (it exists by [4], Theorem 5.1.12, (i) and (iii)), let \mathcal{C} be any locally finite open refinement of \mathcal{B} (it exists by the definition of paracompactness), and let B_p be the intersection of A_p with an arbitrary element of \mathcal{C} containing p . Notice that the B_p form an open cover of $\mathcal{P}(\Omega)$. If Ω is finite, replace $\{B_p\}_{p \in \mathcal{P}(\Omega)}$ by its open shrinking of order $|\Omega| - 1$ (it exists by the Dowker theorem, Theorem 7.2.4 in [4], since Ω is normal, Theorem 3.1.9 in [4], and $\dim(\mathcal{P}(\Omega)) = |\Omega| - 1$, [4], Theorem 7.3.19); we will use the same notation $\{B_p\}_{p \in \mathcal{P}(\Omega)}$ for the shrinking. Let $\{f_s\}_{s \in S}$ be a locally finite partition of unity subordinated to the open cover $\{B_p\}_{p \in \mathcal{P}(\Omega)}$ ([4], Theorem 5.1.9). For each $s \in S$ choose a $p_s \in \mathcal{P}(\Omega)$ such that $\{p \mid f_s(p) > 0\} \subseteq B_{p_s}$. Set, for $\omega \in \Omega$ and $p \in \mathcal{P}(\Omega)$,

$$S^*(\omega, p) := \sum_{s \in S} S_n(\omega, p_s) f_s(p)$$

(notice that only a finite number of addends are non-zero, so the sum is well-defined).

In the previous section we were considering Sceptic's moves $S_n(\omega, p)$ lower semicontinuous in p and satisfying $\int_{\Omega} S_n(\omega, p)p(d\omega) \leq 0$ for all $p \in \mathcal{P}^{\text{fin}}(\Omega)$. It is clear that $S^*(\omega, p)$ is even continuous in p ; let us check that it almost satisfies $\int_{\Omega} S^*(\omega, p)p(d\omega) \leq 0$ for all $p \in \mathcal{P}(\Omega)$. We have:

$$\begin{aligned}
\int_{\Omega} S^*(\omega, p) p(d\omega) &= \int_{\Omega} \sum_{s \in S_p} S_n(\omega, p_s) f_s(p) p(d\omega) \\
&= \sum_{s \in S_p} \int_{\Omega} S_n(\omega, p_s) p(d\omega) f_s(p) \leq \sum_{s \in S_p} \delta f_s(p) = \delta, \quad (3)
\end{aligned}$$

where S_p is the finite set of all s for which $f_s(p) > 0$; the inequality in (3) uses the fact that $p \in B_{p_s} \subseteq A_{p_s}$ and the definition (2). Therefore, $\int_{\Omega} S(\omega, p) p(d\omega) \leq 0$ for all p , where $S := S^* - \delta$. Applying to S the argument given in the proof of Theorem 1, we can see that there exists $p^* \in \mathcal{P}(\Omega)$ satisfying $S(\omega, p^*) \leq 0$, i.e., $S^*(\omega, p^*) \leq \delta$, for all $\omega \in \Omega$.

Make Forecaster select P_n concentrated on the p_s with positive $f_s(p^*)$ and assigning weight $f_s(p^*)$ to each of these p_s . This will ensure that P_n is concentrated on a finite subset, $D(P_n)$, of an element of \mathcal{A}_n and that $|D(P_n)| \leq |\Omega|$.

The rest of the proof proceeds similarly to the proof of Theorem 3 in [15]. Let δ be $\epsilon 2^{-n}$ or less. This will ensure

$$\int S_n(\omega, p) P_n(dp) \leq \epsilon 2^{-n} \quad (4)$$

for all $\omega \in \Omega$. Let Forecaster's strategy further tell him to use as his second move the function f_n given by

$$f_n(p) := \frac{1}{1 + \epsilon} (S_n(\omega_n, p) - \epsilon 2^{-n}) \quad (5)$$

for $p \in D(P_n)$ and defined arbitrarily for $p \notin D(P_n)$. The condition $\int f_n dP_n \leq 0$ is then guaranteed by (4).

It remains to check $\mathcal{K}_n \leq (1 + \epsilon)\mathcal{F}_n$ (this will also establish that \mathcal{F}_n is never negative). This can be done by a formal calculation (as in the proof of Theorem 3 in [15]), but I prefer the following intuitive picture. We would like Forecaster to use $f_n(p) := S_n(\omega_n, p) - \epsilon 2^{-n}$ (for $p \in D(P_n)$) as his second move; this would always keep his capital \mathcal{F}_n above $\mathcal{K}_n - \epsilon$. To make sure that \mathcal{F}_n is never negative, Forecaster would have to start with initial capital $\mathcal{F}_0 = 1 + \epsilon$, which, moreover, would lead to $\mathcal{F}_n \geq \mathcal{K}_n$, $\forall n$; our protocol, however, requires $\mathcal{F}_0 = 1$. Therefore, Forecaster's strategy has to be scaled down to the initial capital 1, leading to (5); $\mathcal{F}_n \geq \mathcal{K}_n$ becomes $(1 + \epsilon)\mathcal{F}_n \geq \mathcal{K}_n$. (Scaling down a strategy to a smaller initial capital means that the player multiplies the strategy's moves by the same factor as he has multiplied the initial capital, thus assuring that the capital on succeeding rounds is also multiplied by this factor.) ■

Corollary 2 *Forecaster has a winning strategy in the randomized game.*

Proof We are required to show that for every legal strategy \mathcal{S} for Forecaster, we can construct another legal strategy \mathcal{S}^* such that whenever \mathcal{S} 's capital is unbounded, \mathcal{S}^* 's tends to infinity. I will reproduce a simple construction (which I learned from Shen) given in [15], the proof of Theorem 3. (For a more efficient, in certain respects, construction see [12], Lemma 3.1; an even better construction has been recently devised by Vereshchagin and Shen.)

We choose some number larger than 1, say 2. Starting, as the game requires, with initial capital 1 for Forecaster, we have him play \mathcal{S} until its capital exceeds 2. Then he sets aside 1 of this capital and continues with a rescaled version of \mathcal{S} , scaled down to the reduced capital. When the capital again exceeds 2, he again sets aside 1, and so forth. The money set aside, which is part of the capital earned by this strategy, grows without bound. ■

Theorem 2 imposes a condition of continuity on Sceptic's move S_n whereas Theorem 1 only requires lower semicontinuity (in a different argument). A natural question is whether we can relax the former condition. The key point in the proof of Theorem 2 where the continuity of S_n in ω is used is the claim that the set (2) is open. This claim will still be true if S_n is only required to be upper semicontinuous in ω , at least when Ω is a metric compact. We did not pursue this generalization since it can be deduced from Theorem 2 as a corollary (Corollary 3 below).

Let us say that a real-valued function f on Ω is *strongly upper semicontinuous* if there is a monotonic sequence of bounded above real-valued functions $f_1 \geq f_2 \geq \dots$ on Ω that converges to f everywhere. For metric compacts, this requirement coincides with upper semicontinuity ([4], Problems 1.7.15(c) and 3.12.23(g)), but in general it is stronger ([4], Problems 1.7.14(a), 1.7.15(c), and 3.12.23(g)).

Corollary 3 *Theorem 2 will continue to hold if the condition that $S_n(\omega, p)$ be continuous in $\omega \in \Omega$ in the randomized game is relaxed to the condition that $S_n(\omega, p)$ be strongly upper semicontinuous in $\omega \in \Omega$.*

Proof The proof proceeds similarly to the end of the proof of Theorem 2. Let ϵ' be a small positive constant (we will need $(1 + \epsilon')^2 \leq 1 + \epsilon$). Fix, for a moment, a round n of the game. By the monotone convergence theorem and the definition of strong upper semicontinuity, there exists a function $S'_n : \Omega \times \mathcal{P}(\Omega) \rightarrow \mathbb{R}$ such that $S'_n \geq S_n$, $S'_n(\omega, p)$ is continuous in $\omega \in \Omega$, and $\int_{\Omega} S'_n(\omega, p)p(d\omega) \leq \epsilon'2^{-n}$ for all $p \in \mathcal{P}(\Omega)$. Theorem 2 is applicable to Sceptic's move

$$S''_n := \frac{1}{1 + \epsilon'} (S'_n - \epsilon'2^{-n})$$

on round n , for each $n = 1, 2, \dots$, and it asserts the existence of a strategy for Forecaster ensuring

$$\mathcal{K}_n \leq \mathcal{K}'_n \leq (1 + \epsilon')\mathcal{K}''_n \leq (1 + \epsilon')^2 \mathcal{F}_n,$$

where \mathcal{K}' is the capital corresponding to the strategy S' (formally, $\mathcal{K}'_n := 1 + \sum_{i=1}^n S'_i(\omega_i, p_i)$) and \mathcal{K}'' is the capital corresponding to the strategy S'' . ■

Theorem 2 is a general form of Theorem 5 in [15] (that theorem is not part of the journal version). This note is self-contained from the mathematical point of view, but for further motivation behind Theorem 2 the reader is referred to [15].

4 Discussion

This note assumes that the outcome space Ω is a compact. This assumption is not as restrictive as it seems since a wide range of topological spaces have compactifications that are still “nice” topological spaces (cf. [13], the subsection on pp. 4–5). It appears that implications of this fact for prediction with expert advice (see, e.g., [14]) deserve to be explored.

Acknowledgments

This work was motivated by Sasha Shen’s question. It was partially supported by EPSRC through grant EP/F002998/1.

References

- [1] Ravi P. Agarwal, Maria Meehan, and Donal O’Regan. *Fixed Point Theory and Applications*. Cambridge University Press, Cambridge, England, 2001.
- [2] Patrick Billingsley. *Convergence of Probability Measures*. Wiley, New York, 1968.
- [3] Richard M. Dudley. *Real Analysis and Probability*. Cambridge University Press, Cambridge, England, revised edition, 2002.
- [4] Ryszard Engelking. *General Topology*. Heldermann, Berlin, second edition, 1989.
- [5] Dean P. Foster and Rakesh V. Vohra. Asymptotic calibration. *Biometrika*, 85:379–390, 1998.
- [6] Peter Gács. Uniform test of algorithmic randomness over a general space. *Theoretical Computer Science*, 341:91–137, 2005.
- [7] 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.
- [8] Bronisław Knaster, Kazimierz Kuratowski, and Stefan Mazurkiewicz. Ein Beweis des Fixpunktsatzes für n -dimensionale Simplexe. *Fundamenta Mathematicae*, 14:132–137, 1929.
- [9] Leonid A. Levin. Uniform tests of randomness. *Soviet Mathematics Doklady*, 17:337–340, 1976. The Russian original in: Доклады АН СССР 227(1), 1976.
- [10] Alvaro Sandroni. The reproducible properties of correct forecasts. *International Journal of Game Theory*, 32:151–159, 2003.

- [11] Alvaro Sandroni, Rann Smorodinsky, and Rakesh V. Vohra. Calibration with many checking rules. *Mathematics of Operations Research*, 28:141–153, 2003.
- [12] Glenn Shafer and Vladimir Vovk. *Probability and Finance: It's Only a Game!* Wiley, New York, 2001.
- [13] Vladimir Vovk. Competing with stationary prediction strategies. Technical Report [arXiv:cs/0607067](https://arxiv.org/abs/cs/0607067) [cs.LG], [arXiv.org](https://arxiv.org/) e-Print archive, July 2006.
- [14] Vladimir Vovk. Predictions as statements and decisions, The Game-Theoretic Probability and Finance project, <http://probabilityandfinance.com>, Working Paper #17, June 2006.
- [15] 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 (last revised August 2007). Journal version: *Journal of the Royal Statistical Society B*, 67:747–763, 2005.
- [16] 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).

Game-Theoretic Probability and Finance Project Defensive Forecasting Subseries Working Papers

7. *Good randomized sequential probability forecasting is always possible*, by Vladimir Vovk and Glenn Shafer, June 2003 (revised August 2007).
8. *Defensive forecasting*, by Vladimir Vovk, Akimichi Takemura, and Glenn Shafer, September 2004 (revised January 2005).
9. *Experiments with the K29 algorithm*, by Vladimir Vovk, October 2004.
10. *Defensive forecasting for linear protocols*, by Vladimir Vovk, Ilia Nouretdinov, Akimichi Takemura, and Glenn Shafer, February 2005 (revised September 2005).
11. *On-line regression competitive with reproducing kernel Hilbert spaces*, by Vladimir Vovk, November 2005 (revised January 2006).
13. *Non-asymptotic calibration and resolution*, by Vladimir Vovk, November 2004 (revised July 2006).
14. *Competitive on-line learning with a convex loss function*, by Vladimir Vovk, May 2005 (revised September 2005).
16. *Competing with wild prediction rules*, by Vladimir Vovk, December 2005 (revised January 2006).
17. *Predictions as statements and decisions*, by Vladimir Vovk, June 2006.
18. *Leading strategies in competitive on-line prediction*, by Vladimir Vovk, August 2007.
20. *Defensive forecasting for optimal prediction with expert advice*, by Vladimir Vovk, August 2007.
21. *Continuous and randomized defensive forecasting: unified view*, by Vladimir Vovk, August 2007.