Insuring against loss of evidence in game-theoretic probability

A. Philip Dawid[a], Steven de Rooij[a,b], Glenn Shafer[c,d], Alexander Shen[e], Nikolai Vereshchagin[f], Vladimir Vovk[d,*]

[a]*Statistical Laboratory, Centre for Mathematical Sciences, University of Cambridge, UK*
[b]*Centrum Wiskunde & Informatica, Amsterdam, Netherlands*
[c]*Rutgers Business School, Newark, NJ, USA*
[d]*Department of Computer Science, Royal Holloway, University of London, Egham, Surrey, UK*
[e]*Laboratoire d'Informatique Fondamentale, CNRS, Marseille, France*
[f]*Department of Mathematical Logic and the Theory of Algorithms, Moscow State University, Moscow, Russia*

Abstract

Statistical testing can be framed as a repetitive game between two players, Forecaster and Sceptic. On each round, Forecaster sets prices for various gambles, and Sceptic chooses which gambles to make. If Sceptic multiplies by a large factor the capital he puts at risk, he has evidence against Forecaster's ability. His capital at the end of each round is a measure of his evidence against Forecaster so far. This can go up and then back down. If you report the maximum so far instead of the current value, you are exaggerating the evidence against Forecaster. In this article, we show how to remove the exaggeration. Removing it means systematically reducing the maximum in such a way that a rival to Sceptic can always play so as to obtain current evidence as good as Sceptic's reduced maximum. We characterize the functions that can achieve such reductions. Because these functions may impose only modest reductions, we think of our result as a method of insuring against loss of evidence. In the context of an actual market, it is a method of insuring against the loss of what an investor has gained so far.

1. Introduction

In game-theoretic probability (see, e.g., Shafer and Vovk 2001) Sceptic tries to prove Forecaster wrong by gambling against him: the values of Sceptic's capital $\mathcal{K}_n$ measure the changing evidence against Forecaster. We assume that Sceptic's initial capital is $\mathcal{K}_0 = 1$, and that Sceptic is required to ensure that $\mathcal{K}_n \geq 0$ at each time n.

Sceptic can lose as well as gain evidence. At a time n when $\mathcal{K}_n$ is large Forecaster's performance looks poor, but then $\mathcal{K}_i$ for some later time i may be lower and make Forecaster look better. Our result will show that, for a modest cost, Sceptic can avoid losing too much evidence.

Suppose we exaggerate the evidence against Forecaster by considering not the current value $\mathcal{K}_n$ of Sceptic's capital but the greatest value so far:

$$\mathcal{K}_n^* := \max_{i \leq n} \mathcal{K}_i.$$

*Corresponding author. Tel.: +44 1784 443426; fax: +44 1784 439786.
Email address: vovk@cs.rhul.ac.uk (Vladimir Vovk)

Preprint submitted to Statistics and Probability Letters *October 6, 2010*

Continuing research started in Shafer et al. (2010a), we show that there are many functions $F : [1, \infty) \to [0, \infty)$ such that

1. $F(y) \to \infty$ as $y \to \infty$ almost as fast as y, and
2. Sceptic's moves can be modified on-line in such a way that the modified moves lead to capital

$$\mathcal{K}'_n \geq F(\mathcal{K}^*_n), \quad n = 1, 2, \ldots. \tag{1}$$

If we are dissatisfied by the asymptotic character of the first of these two conditions, which does not prevent $\mathcal{K}'_n/\mathcal{K}_n$ from becoming very small for some n, we can compromise by putting a fraction c of the initial capital on Sceptic's original moves and the remaining fraction $1 - c$ on the modified moves, thus obtaining capital $c\mathcal{K}_n + (1 - c)\mathcal{K}'_n$ at each time n. This way Sceptic may sacrifice a fraction $1 - c$ of his capital but gets extra insurance against losing evidence. See Section 3 for details.

As we will show, the set of nondecreasing functions F for which (1) can be achieved can be characterized very simply: it is the set of all nondecreasing F that satisfy

$$\int_1^\infty \frac{F(y)}{y^2} \mathrm{d}y \leq 1. \tag{2}$$

Similar results hold in measure-theoretic probability. One similar measure-theoretic result, for the case where Sceptic's strategy is known in advance, is proven in Shafer et al. (2010a) using a simple method based on Lévy's zero-one law. Lévy's zero-one law generalizes to game-theoretic probability (see Shafer et al. 2010b), but in the present article, where Sceptic's strategy is not necessarily known in advance and Sceptic's moves must be modified on-line, we use an entirely different method of proof, based on the idea of stopping and combining capital processes. This idea has been used previously by various authors, e.g., El-Yaniv et al. (2001, Theorem 1, based on Leonid Levin's personal communication) and Shafer and Vovk (2001, Lemma 3.1). We show that it gives optimal results in the setting of this article.

In Section 4 we explain the meaning of our results in the case where Sceptic represents someone actually trying to make money, not a method for testing forecasts. Suppose Sceptic is a gambler (or an investor) who comes to a casino (stock market) with initial capital 1. On each round, we are allowed to observe how she gambles and then gamble on the same outcome, before observing it, and we want to do so in such a way that our capital will always be at least $F(\mathcal{K}^*)$, where F is a fixed nondecreasing function and $\mathcal{K}^*$ is her maximal capital so far. For which functions F can our goal be achieved? For F satisfying (2).

Alternatively, suppose we have some commodity, such as gold, that we want to sell within a fixed period, say a year. We would like to sell it at the point in time during the year when its price is highest, but of course we never know whether the current price will be exceeded later. If F satisfies (2), we have a strategy that guarantees the price $F(\mathcal{K}^*)$, where $\mathcal{K}^*$ is the highest price over the year. This provides an imperfect alternative to buying a floating lookback put option (see, e.g., Hull 2009, Section 24.8); we get less protection, but we get it for free.

The main idea of the proof can also be explained in these terms. For every threshold u we consider the strategy that stops playing when the current capital reaches (or exceeds) u. This corresponds to the function $F_u(y) := u\,\mathbb{I}_{\{y \geq u\}}$. (If E is some property, $\mathbb{I}_{\{E\}}$ is defined to be 1 if E is satisfied and 0 if not.) Now we can mix these strategies according to some probability measure P on u. It remains to notice that every nondecreasing function F satisfying (2) can be represented as such a mixture: $F(y) = \int F_u(y)P(\mathrm{d}u) = \int_1^y uP(\mathrm{d}u)$.

In this article, we use the standard notation $\mathbb{R}$ for the set of real numbers; the set of natural numbers is $\mathbb{N} := \{1, 2, \ldots\}$. The extended real line $[-\infty, \infty]$ is denoted $\overline{\mathbb{R}}$, and we use the convention $\infty + (-\infty) := \infty$.

2. Calibrating exaggerated evidence

Our prediction protocol involves four players: Forecaster, Sceptic, Rival Sceptic, and Reality.

Protocol 1 Competitive scepticism

$\mathcal{K}_0 := 1$ and $\mathcal{K}'_0 := 1$
for $n = 1, 2, \ldots$ **do**
  Forecaster announces $\mathcal{E}_n \in \mathbf{E}$
  Sceptic announces $f_n \in [0, \infty]^{\mathcal{X}}$ such that $\mathcal{E}_n(f_n) \leq \mathcal{K}_{n-1}$
  Rival Sceptic announces $f'_n \in [0, \infty]^{\mathcal{X}}$ such that $\mathcal{E}_n(f'_n) \leq \mathcal{K}'_{n-1}$
  Reality announces $x_n \in \mathcal{X}$
  $\mathcal{K}_n := f_n(x_n)$ and $\mathcal{K}'_n := f'_n(x_n)$
end for

The parameter of the protocol is a set $\mathcal{X}$, from which Reality chooses her moves; $\mathbf{E}$ is the set of all "outer probability contents" on $\mathcal{X}$ (to be defined momentarily). We always assume that $\mathcal{X}$ contains at least two distinct elements. The reader who is not interested in the most general statement of our result can interpret $\mathbf{E}$ as the set of all expectation functionals $\mathcal{E} : f \mapsto \int f \mathrm{d}P$, P being a probability measure on a fixed σ-algebra on $\mathcal{X}$; in this case Sceptic and Rival Sceptic are required to output functions that are measurable w.r. to that σ-algebra.

In general, an *outer probability content* on $\mathcal{X}$ is a function $\mathcal{E} : \overline{\mathbb{R}}^{\mathcal{X}} \to \overline{\mathbb{R}}$ (where $\overline{\mathbb{R}}^{\mathcal{X}}$ is the set of all functions $f : \mathcal{X} \to \overline{\mathbb{R}}$) that satisfies the following four axioms:

1. If $f, g \in \overline{\mathbb{R}}^{\mathcal{X}}$ and $f \leq g$, then $\mathcal{E}(f) \leq \mathcal{E}(g)$.
2. If $f \in \overline{\mathbb{R}}^{\mathcal{X}}$ and $c \in (0, \infty)$, then $\mathcal{E}(cf) = c\mathcal{E}(f)$.
3. If $f, g \in \overline{\mathbb{R}}^{\mathcal{X}}$, then $\mathcal{E}(f + g) \leq \mathcal{E}(f) + \mathcal{E}(g)$.
4. For each $c \in \mathbb{R}$, $\mathcal{E}(c) = c$, where the c in parentheses is the function in $\overline{\mathbb{R}}^{\mathcal{X}}$ that is identically equal to c.

An axiom of σ-subadditivity on $[0, \infty]^{\mathcal{X}}$ is sometimes added to this list, but we do not need it in this article. (And it is surprising how rarely it is needed in general: see, e.g., Shafer et al. 2010b.) In our terminology we follow Hoffmann-Jørgensen (1987) and Shafer et al. (2010b). Upper previsions studied in the theory of imprecise probabilities (see, e.g., de Cooman and Hermans 2008) are closely related to (but somewhat more restrictive than) outer probability contents.

Protocol 1 describes a perfect-information game in which Sceptic tries to discredit the outer probability contents $\mathcal{E}_n$ issued by Forecaster as a faithful description of Reality's $x_n \in \mathcal{X}$. The players make their moves sequentially in the indicated order. At each step Sceptic and Rival Sceptic choose gambles f_n and f'_n on how x_n is going to come out, and their resulting capitals are $\mathcal{K}_n$ and $\mathcal{K}'_n$, respectively. Discarding capital is allowed, but Sceptic and Rival Sceptic are required to ensure that $\mathcal{K}_n \geq 0$ and $\mathcal{K}'_n \geq 0$, respectively; this is achieved by requiring that f_n and f'_n should be nonnegative.

Let us call a nondecreasing function $F : [1, \infty) \to [0, \infty)$ a *capital calibrator* if there exists a strategy for Rival Sceptic that guarantees (1) with $F(\infty)$ understood to be $\lim_{y \to \infty} F(y)$. We say that a capital calibrator F *dominates* a capital calibrator G if $F(y) \geq G(y)$ for all $y \in [1, \infty)$. We say that F *strictly dominates* G if F dominates G and $F(y) > G(y)$ for some $y \in [1, \infty)$. A capital calibrator is *admissible* if it is not strictly dominated by any other capital calibrator.

Theorem 1. 1. *A nondecreasing function $F : [1, \infty) \to [0, \infty)$ is a capital calibrator if and only if it satisfies (2).*

2. *Any capital calibrator is dominated by an admissible capital calibrator.*
3. *A capital calibrator is admissible if and only if it is right-continuous and*

$$\int_1^\infty \frac{F(y)}{y^2} \mathrm{d}y = 1. \tag{3}$$

Proof. First we prove that any nondecreasing function $F : [1, \infty) \to [0, \infty)$ satisfying

$$F(y) = \int_{[1,y]} u P(\mathrm{d}u), \quad \forall y \in [1, \infty), \tag{4}$$

for a probability measure P on $[1, \infty)$ is a capital calibrator. For each $u \geq 1$, define the following strategy for Rival Sceptic: at step n, the strategy outputs

$$f_n^{(u)} := \begin{cases} f_n & \text{if } \mathcal{K}_{n-1}^* < u \\ u & \text{otherwise} \end{cases}$$

as Rival Sceptic's move f_n'. Let us check that this is a valid strategy, i.e., that $\mathcal{E}_n(f_n^{(u)}) \leq \mathcal{K}_{n-1}^{(u)}$, $n \in \mathbb{N}$, where $\mathcal{K}^{(u)}$ is defined by $\mathcal{K}_0^{(u)} := 1$ and $\mathcal{K}_n^{(u)} := f_n^{(u)}(x_n)$ for $n \in \mathbb{N}$. There are three cases to consider:

1. If $\mathcal{K}_{n-1}^* < u$, we have $\mathcal{K}_{n-1}^{(u)} = \mathcal{K}_{n-1}$ and $\mathcal{E}_n(f_n^{(u)}) = \mathcal{E}_n(f_n) \leq \mathcal{K}_{n-1} = \mathcal{K}_{n-1}^{(u)}$.
2. If n is the smallest number for which $\mathcal{K}_{n-1}^* \geq u$, we have $\mathcal{K}_{n-1}^{(u)} = \mathcal{K}_{n-1} \geq u$ and $\mathcal{E}_n(f_n^{(u)}) = \mathcal{E}_n(u) = u \leq \mathcal{K}_{n-1}^{(u)}$.
3. Otherwise, we have $\mathcal{K}_{n-1}^{(u)} = u$ and so $\mathcal{E}_n(f_n^{(u)}) = \mathcal{E}_n(u) = u = \mathcal{K}_{n-1}^{(u)}$.

Set $f_n'(x) := \int_{[1,\infty)} f_n^{(u)}(x) P(\mathrm{d}u)$, $x \in \mathcal{X}$; this gives $\mathcal{K}_n' = \int_{[1,\infty)} \mathcal{K}_n^{(u)} P(\mathrm{d}u)$ when we set x to x_n. Let us check that this is a valid strategy for Rival Sceptic, i.e., that $\mathcal{E}_n(f_n') \leq \mathcal{K}_{n-1}'$ for all $n \in \mathbb{N}$. This is now obvious if $\mathcal{E}_n$ are expectation functionals, and in general we have

$$\begin{aligned}
\mathcal{E}_n(f_n') &= \mathcal{E}_n\left(\int_{[1,\infty)} f_n^{(u)} P(\mathrm{d}u)\right) \\
&= \mathcal{E}_n\left(\int_{[1,\infty)} \left(\mathbb{I}_{\{\mathcal{K}_{n-1}^* < u\}} f_n + \mathbb{I}_{\{\mathcal{K}_{n-1}^* \geq u\}} u\right) P(\mathrm{d}u)\right) \\
&= \mathcal{E}_n\left(P((\mathcal{K}_{n-1}^*, \infty)) f_n + \int_{[1,\mathcal{K}_{n-1}^*]} u P(\mathrm{d}u)\right) \\
&\leq P((\mathcal{K}_{n-1}^*, \infty)) \mathcal{K}_{n-1} + \int_{[1,\mathcal{K}_{n-1}^*]} u P(\mathrm{d}u) \\
&= \int_{(\mathcal{K}_{n-1}^*,\infty)} \mathcal{K}_{n-1} P(\mathrm{d}u) + \int_{(\mathcal{K}_{n-2}^*,\mathcal{K}_{n-1}^*]} u P(\mathrm{d}u) + \int_{[1,\mathcal{K}_{n-2}^*]} u P(\mathrm{d}u)
\end{aligned}$$

$$\leq \int_{(\mathcal{K}^*_{n-1},\infty)} \mathcal{K}^{(u)}_{n-1} P(\mathrm{d}u) + \int_{(\mathcal{K}^*_{n-2},\mathcal{K}^*_{n-1}]} \mathcal{K}^{(u)}_{n-1} P(\mathrm{d}u) + \int_{[1,\mathcal{K}^*_{n-2}]} \mathcal{K}^{(u)}_{n-1} P(\mathrm{d}u)$$
$$= \int_{[1,\infty)} \mathcal{K}^{(u)}_{n-1} P(\mathrm{d}u) = \mathcal{K}'_{n-1}.$$

The last inequality used the analysis of the three cases above. For small values of n, our convention was $\mathcal{K}^*_0 := 1$ and $\mathcal{K}^*_{-1} := 1$. Notice that our argument only used Axioms 2–4 for outer probability contents; no σ-subadditivity was required. This strategy will guarantee

$$\mathcal{K}'_n = \int_{[1,\infty)} \mathcal{K}^{(u)}_n P(\mathrm{d}u) \geq \int_{[1,\mathcal{K}^*_n]} \mathcal{K}^{(u)}_n P(\mathrm{d}u) \geq \int_{[1,\mathcal{K}^*_n]} u P(\mathrm{d}u) = F(\mathcal{K}^*_n). \tag{5}$$

We can now finish the proof of the statement "if" in part 1 of the theorem, which says that any nondecreasing function $F : [1,\infty) \to [0,\infty)$ satisfying (2) is a capital calibrator. Without loss of generality we can assume that F is right-continuous and that (3) holds. It remains to apply Lemma 1 below.

Let us now check that every capital calibrator satisfies (2). Suppose a capital calibrator F violates (2). We can decrease F so that, for some $a > 1$ and $N \in \mathbb{N}$, it is constant in each interval $[a^{n-1}, a^n)$, $n = 1, \dots, N$, is zero in $[a^N, \infty)$, and still violates (2). Of course, F is still a capital calibrator. The substitution $x = 1/y$ shows that $\int_0^1 F(1/x)\mathrm{d}x > 1$, which can be rewritten as

$$F(1)\left(1 - \frac{1}{a}\right) + F(a)\left(\frac{1}{a} - \frac{1}{a^2}\right) + \cdots + F(a^{N-1})\left(\frac{1}{a^{N-1}} - \frac{1}{a^N}\right) > 1. \tag{6}$$

Suppose, without loss of generality, that $\mathcal{X} \supseteq \{0,1\}$, and let Forecaster always choose

$$\mathcal{E}_n(f) := \frac{1}{a} f(1) + \left(1 - \frac{1}{a}\right) f(0), \quad n \in \mathbb{N}.$$

Let Sceptic play the strategy of always betting all his capital on 1: $f_n(1) := a\mathcal{K}_{n-1}$ and $f_n(x) := 0$ for $x \neq 1$. Then $\mathcal{K}^*_N = a^n$ where n is the number of 1s output by Reality before the first element different from 1 (except that $n = N$ if Reality outputs only 1s during the first N steps). Backward induction shows that the initial capital $\mathcal{K}'_0$ required to ensure $\mathcal{K}'_N \geq F(\mathcal{K}^*_N)$ must be at least

$$F(a^N)\left(\frac{1}{a}\right)^N + F(a^{N-1})\left(\frac{1}{a}\right)^{N-1}\left(1 - \frac{1}{a}\right) + F(a^{N-2})\left(\frac{1}{a}\right)^{N-2}\left(1 - \frac{1}{a}\right)$$
$$+ \cdots + F(a)\frac{1}{a}\left(1 - \frac{1}{a}\right) + F(1)\left(1 - \frac{1}{a}\right) > 1;$$

the inequality follows from (6), but we know that it is false as $\mathcal{K}'_0 = 1$.

We have proved part 1 of the theorem. Part 3 is now obvious, and part 2 follows from parts 1 and 3. □

The following lemma was used in the proof of Theorem 1.

Lemma 1. *A nondecreasing right-continuous function $F : [1,\infty) \to [0,\infty)$ satisfies (3) if and only if (4) holds for some probability measure P on $[1,\infty)$.*

Proof. Let us first check that the existence of a probability measure P satisfying (4) implies (3). We have:

$$\int_{[1,\infty)} \frac{F(y)}{y^2}\mathrm{d}y = \int_{[1,\infty)} \int_{[1,y]} \frac{u}{y^2} P(\mathrm{d}u)\mathrm{d}y = \int_{[1,\infty)} \int_{[u,\infty)} \frac{u}{y^2}\mathrm{d}y P(\mathrm{d}u) = \int_{[1,\infty)} P(\mathrm{d}u) = 1. \quad (7)$$

It remains to check that any nondecreasing right-continuous $F : [1,\infty) \to [0,\infty)$ satisfying (3) satisfies (4) for some probability measure P on $[1,\infty)$. Let Q be the measure on $[1,\infty)$ (σ-finite but not necessarily a probability measure) with distribution function F, in the sense that $Q([1,y]) = F(y)$ for all $y \in [1,\infty)$. Set $P(\mathrm{d}u) := (1/u)Q(\mathrm{d}u)$. We then have (4), and the calculation (7) shows that the σ-finite measure P must be a probability measure (were it not, we would not have an equality in (3)). ☐

According to (3), the functions

$$F(y) := \alpha y^{1-\alpha} \quad (8)$$

are admissible capital calibrators for any $\alpha \in (0,1)$.

3. Insuring against loss of evidence

Condition (2) implies $\liminf_{y\to\infty} F(y)/y = 0$. Therefore, as we mentioned in Section 1, it is possible that $\mathcal{K}'_n/\mathcal{K}_n$ will be very small for some n, and we pointed out a simple way to use Theorem 1 for insuring against this possibility. The following corollary says that it leads to an optimal result.

Corollary 1. *Let $c \ge 0$ and $F : [1,\infty) \to [0,\infty)$ be a nondecreasing function. Rival Sceptic has a strategy ensuring*

$$\mathcal{K}'_n \ge c\mathcal{K}_n + F(\mathcal{K}^*_n) \quad (9)$$

if and only if c and F satisfy

$$\int_1^\infty \frac{F(y)}{y^2}\mathrm{d}y \le 1 - c. \quad (10)$$

Proof. Suppose (10) is satisfied; in particular, $c \in [0,1]$. Using $cf_n + (1-c)f'_n$ as Rival Sceptic's strategy, where f_n are Sceptic's moves and f'_n are Rival Sceptic's moves guaranteeing $\mathcal{K}'_n \ge \frac{1}{1-c}F(\mathcal{K}^*_n)$ (cf. Theorem 1), we can see that Rival Sceptic can guarantee (9).

Now suppose Rival Sceptic can ensure (9), but (10) is violated. As in the proof of Theorem 1, we can decrease F so that, for some $a > 1$ and $N \in \mathbb{N}$, it is constant in each interval $[a^{n-1}, a^n)$, $n = 1, \dots, N$, is zero in $[a^N, \infty)$, and still violates (10). Similarly to (6), we have

$$F(1)\left(1 - \frac{1}{a}\right) + F(a)\left(\frac{1}{a} - \frac{1}{a^2}\right) + \cdots + F(a^{N-1})\left(\frac{1}{a^{N-1}} - \frac{1}{a^N}\right) > 1 - c.$$

Suppose $\mathcal{X} \supseteq \{0,1\}$ and define Forecaster's and Sceptic's strategies as before. Now backward induction shows that the initial capital $\mathcal{K}'_0$ required to ensure $\mathcal{K}'_N \ge c\mathcal{K}_N + F(\mathcal{K}^*_N)$ must be at least

$$ca^N\left(\frac{1}{a}\right)^N + F(a^N)\left(\frac{1}{a}\right)^N + F(a^{N-1})\left(\frac{1}{a}\right)^{N-1}\left(1 - \frac{1}{a}\right)$$
$$+ F(a^{N-2})\left(\frac{1}{a}\right)^{N-2}\left(1 - \frac{1}{a}\right) + \cdots + F(a)\frac{1}{a}\left(1 - \frac{1}{a}\right) + F(1)\left(1 - \frac{1}{a}\right)$$

$$> c + (1-c) = 1.$$

This contradicts $\mathcal{K}'_0 = 1$. □

According to (8) and (9), Rival Sceptic can guarantee

$$\mathcal{K}'_n \geq c\mathcal{K}_n + (1-c)\alpha(\mathcal{K}^*_n)^{1-\alpha} \tag{11}$$

for any constants $c \in [0,1]$ and $\alpha \in (0,1)$.

Corollary 1 does not mean that (11) or, more generally, (9) cannot be improved; it only says that the improvement will not be significant enough to decrease the coefficient in front of $\mathcal{K}_n$. For example, if we do not discard the term $\int_{(\mathcal{K}^*_n,\infty)} \mathcal{K}^{(u)}_n P(\mathrm{d}u)$ in (5), we will obtain

$$\mathcal{K}'_n \geq P((\mathcal{K}^*_n, \infty))\mathcal{K}_n + F(\mathcal{K}^*_n). \tag{12}$$

The coefficient $P((\mathcal{K}^*_n, \infty))$ in front of $\mathcal{K}_n$ tends to zero as $\mathcal{K}^*_n \to \infty$.

In particular, using (12) allows us to improve (11) to

$$\mathcal{K}'_n \geq c\mathcal{K}_n + (1-c)(1-\alpha)(\mathcal{K}^*_n)^{-\alpha}\mathcal{K}_n + (1-c)\alpha(\mathcal{K}^*_n)^{1-\alpha}.$$

4. Insuring against loss of money

In conclusion, we discuss an application of our results to finance. Consider a financial market in which K securities are traded over successive periods. Recall that the *return* of a security during a trading period is the ratio

$$\frac{\text{closing price} - \text{opening price}}{\text{opening price}},$$

and let $x_{n,k}$ be the kth security's return in the nth trading period. For each period n, write x_n for the vector $(x_{n,1}, \ldots, x_{n,K})$, which is in $\mathcal{X} := [-1, \infty)^K$.

Now consider how an investor might invest in the market during period n. Write $\gamma_{n,k}$ for the amount of money invested in security k during period n, and write γ_n for the vector $(\gamma_{n,1}, \ldots, \gamma_{n,K})$. Under the simplifying assumption that the investor is allowed to go long or short by any amount, γ_n can be any vector in $\mathbb{R}^K$. If the investor chooses γ_n and the market chooses x_n, then the investor's profit will be $\gamma_{n,1}x_{n,1} + \cdots + \gamma_{n,K}x_{n,K}$.

This simple model of a financial market can be embedded in Protocol 1 as follows. As we said, $\mathcal{X} := [-1, \infty)^K$. At each step Forecaster chooses the same outer probability content $\mathcal{E}_n = \mathcal{E}$ on $\mathcal{X}$, which is defined by

$$\mathcal{E}(f) := \inf\left\{\mathcal{K} \mid \exists\gamma \in \mathbb{R}^K \forall x \in \mathcal{X} : \mathcal{K} + \gamma_1 x_1 + \cdots + \gamma_K x_K \geq f(x)\right\}.$$

We leave it to the reader to verify that this satisfies the axioms for an outer probability content. In the situation of Protocol 1, where the function f is nonnegative, the infimum does not change if we additionally require that $\gamma_1, \ldots, \gamma_K$ should be nonnegative and sum to at most $\mathcal{K}$, and therefore, inf is attained and can be replaced with min.

Now Forecaster is a dummy player, Sceptic is an investor in the market, Rival Sceptic is another investor, who decides on his own investment for each trading period after seeing Sceptic's decision, and Reality is the market. Results of this article show that Rival Sceptic can modify

Sceptic's decisions in such a way that his capital $\mathcal{K}'_n$ never drops much below the maximal value $\mathcal{K}^*_n$ achieved by Sceptic's capital $\mathcal{K}_i$ so far. For example, for any constants $c \in [0,1]$ and $\alpha \in (0,1)$, Rival Sceptic can guarantee (11).

We can apply these ideas not only to securities but also to commodities or dynamic portfolios of securities. In particular, our two stories at the end of Section 1 are special cases of the framework of this section corresponding to $K = 1$. (The case of an arbitrary K is not really more general: as far as our results are concerned, it reduces to the case of $K = 1$.)

Acknowledgements

Thanks to Peter Grünwald for inspiring part of this work. Comments by a referee were very helpful; in particular, they led to the addition of Section 4 and expansion of Section 1. This research has been supported in part by ANR grant NAFIT ANR-08-EMER-008-01 and EPSRC grant EP/F002998/1.

References

de Cooman, G., Hermans, F., 2008. Imprecise probability trees: bridging two theories of imprecise probability. Artificial Intelligence 172, 1400–1427.

El-Yaniv, R., Fiat, A., Karp, R. M., Turpin, G., 2001. Optimal search and one-way trading online algorithms. Algorithmica 30, 101–139.

Hoffmann-Jørgensen, J., 1987. The general marginal problem. In: Kurepa, S., Kraljević, H., Butković, D. (Eds.), Functional Analysis II. Vol. 1242 of Lecture Notes in Mathematics. Springer, Berlin, pp. 77–367.

Hull, J. C., 2009. Options, Futures, and other Derivatives, 7th Edition. Pearson Prentice-Hall, Upper Saddle River, NJ.

Shafer, G., Shen, A., Vereshchagin, N., Vovk, V., 2010a. Test martingales, Bayes factors, and p-values. The Game-Theoretic Probability and Finance project, Working Paper 33, `http://probabilityandfinance.com`.

Shafer, G., Vovk, V., 2001. Probability and Finance: It's Only a Game! Wiley, New York.

Shafer, G., Vovk, V., Takemura, A., 2010b. Lévy's zero-one law in game-theoretic probability. The Game-Theoretic Probability and Finance project, Working Paper 29, `http://probabilityandfinance.com`.