On a Review by V. N. Tutubalin

Glenn Shafer
Rutgers School of Business - Newark and New Brunswick
180 University Avenue, Newark, New Jersey 07102 USA
gshafer@andromeda.rutgers.edu
http://www.andromeda.rutgers.edu/~gshafer/

Vladimir Vovk
Department of Computer Science
Royal Holloway, University of London
Egham, Surrey TW20 0EX UK
vovk@cs.rhul.ac.uk

March 28, 2002

Valery Nikolaevich Tutubalin, of the Department of Probability Theory at Moscow University, has recently reviewed our book, *Probability and Finance: It's Only a Game!* (Wiley, 2001). It is our understanding that his review will appear in Russian in "Теория вероятностей и ее применения", the Russian journal that is regularly translated into English as *Theory of Probability and Its Applications*.

Professor Tutubalin raises important questions about the range of usefulness of our book's game-theoretic framework for probability. In this response, we discuss his questions in some detail.

We have translated the review into English, and with Professor Tutubalin's permission we have posted both the Russian text and our translation at the web site for our book, http://www.cs.rhul.ac.uk/~vovk/book/. This response can also be downloaded from there.

1 Introduction

We are pleased that Professor Tutubalin has reviewed our book, and we are doubly delighted that he thinks so highly of it. We hope that his perceptive, informative, and highly complimentary review will be widely read.

Tutubalin takes issue with us on one important point. Although he feels that our game-theoretic conception of probability is perfectly adapted to the study of financial markets, where probability games are actually played and money changes hands, he is not convinced that this conception fits other applications of mathematical probability. So he issues a challenge: "Начну некоторую полемику с того утверждения, что, скорее всего, вряд ли возможна такая концепция вероятности, которая удовлетворит все возможные приложения." In our translation: "I will start a controversy by asserting that a single conception of probability suitable for all possible applications may scarcely be possible." We are delighted to take up this offer of further discussion, and we anticipate a more amicable discussion than the words "полемика" and "controversy" might suggest, for we find much with which to agree in Tutubalin's position. We completely agree, in fact, that there is more than one useful conception of probability, and that no single conception is suitable for all applications. We do think that Tutubalin is mistaken to restrict the game-theoretic conception to situations where money actually changes hands (see the quote in §3).

The starting point of our framework for probability is a story—a story about a game where money does change hands. We may call it *the probability story*. The key to appreciating how widely the probability story can be used is to step back and think in general about the different ways any story can be used. We can use a story as a straightforward description of what is happening in front of our eyes, but we can also use it in other ways. We can compare what happens in front of our eyes with the story—we can use the story as a standard against which to compare what we see. We can also try to understand what we see in front of our eyes by studying ways it can fit into the story. When we use the probability story as a straightforward description of a something happening in front of our eyes, money must be changing hands in front of our eyes. But we can use the story in other ways in situations where money does not change hands.

In [2], Shafer elaborates on this point in order to show how broad the game-theoretic framework for probability really is. Here we do not dwell further on the general point, but we discuss three specific examples where the probability story is used without money actually changing hands: probabilistic weather forecasting (§2), quantum mechanics (§4), and the theory of errors (§3). We conclude, in §5, with comments on another issue that Tutubalin raises: the scaling of the efficient market hypothesis.

2 Evaluating a Probability Forecaster

The discussion of probability forecasting in our book is spread over several different passages (pp. 7–8, 57–58, 162–164, and 177–182). We will draw together the basic points before quoting and responding to Tutubalin's comments.

2.1 The game-theoretic treatment

Suppose a weather forecaster gives probabilities each day for some event, say rain, on the following day.

Players: Forecaster, Reality
Protocol:
  FOR $n = 1, 2, \dots, N$:
    Forecaster announces $p_n \in [0, 1]$.
    Reality announces $x_n \in \{0, 1\}$.

On each round, Forecaster gives a probability for rain (p_n), and then Reality decides whether it rains ($x_n = 1$) or not ($x_n = 0$).

What do Forecaster's p_n mean? Forecaster is not literally offering to bet, but the p_n are most often understood in terms of betting. We can imagine Forecaster issuing this challenge to his listeners:

> I don't have the time or money to make bets with you, but if I did really offer you the opportunity to bet with me at the odds I set, you couldn't make a lot of money from me.

To make sense of this, we can introduce an imaginary third player, Skeptic, who is allowed to bet:

Players: Forecaster, Skeptic, Reality
Protocol:
  $\mathcal{K}_0 := 1$.
  FOR $n = 1, 2, \dots, N$:

Forecaster announces $p_n \in [0,1]$.
Skeptic announces $M_n \in \mathbb{R}$.
Reality announces $x_n \in \{0,1\}$.
$\mathcal{K}_n := \mathcal{K}_{n-1} + M_n(x_n - p_n)$.

If Skeptic chooses M_n positive, he is betting on rain at odds $p_n : (1 - p_n)$ (he loses $M_n p_n$ if it does not rain but gains $M_n(1 - p_n)$ if it does rain). If he chooses M_n negative, he is betting against rain at these same odds.

We make more precise the idea that an opponent cannot make a lot of money betting against the weather forecaster by adopting Cournot's principle for this protocol: if Skeptic does not risk bankruptcy, then he will not achieve a very large final capital $\mathcal{K}_N$. This is our way of giving meaning to the p_n. There may be other ways of interpreting these numbers, but we do not know of any that are equally clear and precise. And because Cournot's principle can also be used to give a very natural interpretation of the classical limit theorems of probability theory (this is the main topic of the first half of our book), we can argue that our way of interpreting the p_n is fully in the spirit of classical probability.

This protocol gives an example where the probability story is used indirectly: no real money changes hands, and Skeptic's betting in imaginary money is nothing more than a means of testing Forecaster's claim.

Once we adopt Cournot's principle, testing Forecaster's probabilities becomes a matter of trying out strategies for Skeptic. Because Skeptic makes his move M_n just after Forecaster makes his move p_n, Skeptic can use any information that is available in the world at this point in order to choose M_n. If we find a rule, a computer program, or even a rival weather expert whose choices of the M_n using such information produces a large $\mathcal{K}_N$, then we have empirically refuted Forecaster's probabilities.

A weather forecaster may, of course, make only a more modest claim for his probabilities. Instead of defying all comers to beat his odds, he might limit his challenge to people who have no more information than he has. If he is working in the United States, for example, he might claim that no one can beat his odds working only with information supplied by the National Weather Service. In this case, we would test his claim using only strategies based on this information.

In Chapter 7 of our book, we show that the strategy that establishes Lindeberg's central limit theorem can be used to test Forecaster's p_n. In this context, the number of rounds N depends on Forecaster: we continue the game until $\sum_{n=1}^{N} p_n(1-p_n)$ is sufficiently large. The strategy that establishes Lindeberg's theorem will produce a large final capital $\mathcal{K}_N$ for Skeptic, thus refuting Forecaster's p_n, if

$$\frac{\sum_{n=1}^{N}(x_n - p_n)}{\sqrt{\sum_{n=1}^{N} p_n(1-p_n)}} \tag{1}$$

has a value that would be extreme for the standard normal distribution. So if the absolute value of (1) exceeds 3, say, we may argue that the p_n have been falsified. As we mention in the book, this procedure tests a particular aspect of the p_n—their overall calibration. It rejects the p_n when their average, $\sum p_n/N$, is too different from the empirical frequency of rain, $\sum x_n/N$. In this case, Skeptic can make money from Forecaster based on very little knowledge indeed; all he needs to know is the previous moves in the game.

2.2 Tutubalin's comments on weather forecasting

Tutubalin writes as follows (in our translation) in reference to our test of calibration:

> In the matter of weather forecasting, the authors propose to test the soundness of a calculation of probabilities for rain on successive days using a statistic that can also be obtained from the frequentist conception of probability, normalized as if the differences $x_i - p_i$ were statistically independent, where x_i is the indicator of the event "presence of rain on day i", and p_i is the forecasted probability of rain. Since nothing guarantees the independence of such differences, I would also recommend thinning out the sequence of observations in different ways, so as to see how the statistic comes out just on Mondays, just on Tuesdays, etc., and then comparing these values with the common value of the statistic. Who will convince me that such an attempt to weaken the assumed statistical dependence[1] is meaningless? Yet it is not, for some reason, recommended under the game-theoretic approach.

[1] Perhaps our translation should be improved here. Apparently, what Tutubalin proposes to weaken is an assumption *about* statistical dependence, namely, the assumption that the errors in the forecasts are statistically *in*dependent of each other or of the day of the week.

This passage is misleading, we feel, in several respects.

Most importantly, contrary to the last sentence of the quotation, we *do* recommend the additional tests recommended by Tutubalin and many other additional tests as well. Overall calibration is only one of many aspects of the p_n that can be tested. We regret that we did not make this point more clearly and emphatically in Chapter 7.

Tutubalin's comments about statistical independence are also potentially misleading. The first source of confusion is that several different concepts of independence and dependence seem to be in play in the quoted passage. This is true even if we take a purely measure-theoretic point of view—i.e., if we assume that the $x_n - p_n$ have a joint probability distribution. First, we can discuss independence of the $x_n - p_n$ of each other, in the sense that their joint distribution equals the product of their marginal distributions. Second, we can talk about their independence of the day of the week; one interpretation of this is that their marginal distributions coincide. Let us say that the $x_n - p_n$ are *identically distributed* when the second notion is meant, and that the $x_n - p_n$ are *iid* when both are meant.

We do not assume that the differences $x_n - p_n$ are independent in any of these senses. Tutubalin seems at first to acknowledge this; he asserts only that the result we derive "can also be obtained" from an argument relying on statistical independence (presumably he has the iid assumption in mind here). But by the end of the passage, the assumption that the $x_n - p_n$ are independent of the day of the week appears as something belonging to our approach, which needs to be weakened.[2]

In our story, Skeptic is entitled to use any information he has when Forecaster announces p_n, including his information about the day of the week, in deciding on his move M_n. He can decide that he will bet only on Mondays, and hence the game-theoretic argument that leads to the asymptotic normality of (1) when this statistic is calculated over all days also applies when it is

[2]The passage also seems to suggest that one needs the iid assumption, or at least mutual independence of the $x_n - p_n$, in order to derive the asymptotic normality of (1) within the measure-theoretic framework, and this too is misleading. The asymptotic normality of (1) can indeed be derived assuming that the x_n are independent of each other and that p_n is the expected value of x_n, and under these assumptions the $x_n - p_n$ are certainly independent of each other. But it can also be derived assuming only that p_n is the conditional expected value of x_n given $x_1, \ldots, x_{n-1}$, with no assumption of independence. Under this assumption, the p_n are random, and the $x_n - p_n$ are uncorrelated with each other but not independent of each other.

calculated only for Mondays. It also applies when the statistic is calculated only for days with some other property that occurs sufficiently often and is known by the time Forecaster announces p_n. In this sense, the statistical behavior of $\sum(x_n - p_n)/\sqrt{\sum p_n(1-p_n)}$ should be independent of whether it is calculated for all days, only for Mondays, or only for days with some other property. This is a kind of independence (a third notion of independence) of the $x_n - p_n$ from everything else. But it is a consequence of our adoption of Cournot's principle, not an assumption we make directly. And it does not involve the $x_n - p_n$ being stochastically independent, in any measure-theoretic or frequentist sense, from each other or from anything else.[3]

3 The Theory of Errors

In the example of the weather forecaster, Forecaster and Reality were playing in front of our eyes, and we invented a third player, Skeptic. Now we consider an example where Forecaster is hidden from us.

3.1 Averaging out the errors

Consider, for simplicity, the case of repeated measurements of an unknown physical quantity θ, such as the distance to the sun. Suppose the error of each measurement is bounded in absolute value by a known quantity C, and suppose further that we consider the measurements unbiased, in the sense that a person who is given the opportunity to buy and sell the errors (which may turn out positive or negative) at price zero would not be able to multiply his initial stake substantially without risking bankruptcy. Writing $y_1, \ldots, y_N$ for the measurements, we can describe this situation with the following game:

Players: Forecaster, Skeptic, Reality
Parameters: N, $C > 0$, $\alpha > 0$, $\epsilon > 0$
Protocol:
  Forecaster announces $\theta \in \mathbb{R}$.
  $\mathcal{K}_0 := \alpha$.
  FOR $n = 1, 2, \ldots, N$:

[3]Professor Tutubalin has informed us that the only notion of independence used in the quoted passage is the standard statistical independence (the one implying that the variance of the total number of successes in a series of trials equals the sum of variances in the individual trials).

Skeptic announces $M_n \in \mathbb{R}$.
Reality announces $y_n \in [\theta - C, \theta + C]$.
$\mathcal{K}_n := \mathcal{K}_{n-1} + M_n(y_n - \theta)$.

Rule for Winning: Skeptic wins if $\mathcal{K}_n$ is never negative and either $\mathcal{K}_N \geq 1$ or $|\overline{y} - \theta| < \epsilon$, where $\overline{y} := \sum_{n=1}^N y_n$.

This game is a variant of the game for the weak law of large numbers on p. 124 of our book. Proposition 6.1 on p. 125, adapted to this variant, tells us that Skeptic has a winning strategy if $N \geq C^2/(\alpha\epsilon)^2$.

To use this result, we choose the initial stake α and the error bound ϵ small. Because α is small, the result $\mathcal{K}_N \geq 1$ means that Skeptic has multiplied his initial stake by the large factor $1/\alpha$, and this is ruled out if we adopt Cournot's principle. So the fact that Skeptic has a winning strategy, together with Cournot's principle, tells us that we can expect $|\overline{y} - \theta|$ to be small when N is sufficiently large. In other words, the errors average out, and the average measurement $\overline{y}$ will be close to θ.

It should be kept in mind that this is a perfect information game; Skeptic sees the move θ even though we, as outsiders to the game, do not see it. What is important is that even we outsiders have all information to compute the estimate $\overline{y}$ of θ.

3.2 Tutubalin on the theory of errors

Tutubalin expresses his skepticism about the ability of our framework to handle errors of measurement as follows:

> ... Here the devil's move x_n is not a change in the price of an asset but rather an error in the measurement of a physical quantity in some experiment. It would be very good to be able to study conditions under which the law of large numbers is satisfied, i.e., the error disappears as the result of averaging a great number of observations. But who is in a position to make monetary bets on the errors of experiments? The game-theoretic conception clearly has nothing to do with this situation.

The key part of this quotation is the question: "... who is in a position to make monetary bets on the errors of experiments?" If we write x_n for the error, as Tutubalin suggests, then $x_n = y_n - \theta$. We observe the measurement y_n when it is made, but we do not know θ, and therefore the error x_n remains unknown to us. So Tutubalin's assertion that no one is in a position to bet on x_n is correct: no one can bet on x_n, because as long as θ is unknown, the bet cannot be settled.

In some cases the physical quantity θ is eventually known with enough accuracy that bets on the errors could eventually be settled, but this does not really affect Tutubalin's point, because our game requires that x_n should become known to the bettor (Skeptic) in the course of the game, before he decides how to bet on x_{n+1}.

Our response to Tutubalin is rather that Skeptic does not need to be a real player here, just as he does not need to be a real player in the case of weather forecasting. Skeptic is theoretical, and so he may see things that we do not see. The quantity θ is theoretical, in the sense that we do not know its value. If it makes sense to reason about this theoretical quantity, then it makes sense to consider the theoretical hypothesis that someone who knows its value cannot multiply his capital substantially without risking bankruptcy.

4 Quantum Mechanics

In our book (pp. 189–191), we explained how the game-theoretic framework accommodates one version of John von Neumann's axioms for quantum mechanics. Tutubalin appears to feel that this venture into physics was overly ambitious. Probably we should have emphasized more that our remarks have very little to do with most discussions of the philosophical foundations of quantum mechanics (see, e.g., [1]); we are only concerned with its well-established core. Our goal was to attract the reader's attention to the fact that measure-theoretic probability has to be stretched uncomfortably to accommodate this core.

4.1 A quantum-mechanical law of large numbers

Consider, for example, the following law of large numbers, which is the only statement we prove in this context (p. 191 of the book):

Corollary 8.5 *Suppose Observer repeatedly measures observables bounded by some constant C. Skeptic can force the event*

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \left(a_n(i_n) - \int a_n dp_n \right) = 0. \tag{2}$$

How can we express this statement in terms of measure-theoretic probability? The only way we can see is to define some underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and random elements a_n, i_n, and p_n and to assert that the event (2) happens $\mathbb{P}$-almost surely. This can be done, of course, but just think what an unattractive object $\mathbb{P}$ would be. This probability measure will have to answer very different kinds of questions, such as:

- What is the probability that the observer will observe his particle (with wave function ϕ) in the region V?
- What is the probability that the observer will decide to measure the position of his particle, rather than its momentum?

Probabilities of the first kind are uncontroversial; indeed, they are supplied by quantum mechanics itself. Probabilities of the second kind, even if they exist, cannot be derived from the primitive version of quantum mechanics (Axioms 8.1–8.5) that we consider. One might hope to derive it from the observer's psychology, or maybe from a more sophisticated theory of everything, but why should it be necessary to go so far afield?

The fact that our game-theoretic theoretic framework permits a quantum-mechanical law of large numbers without any appeal to the existence of probabilities other than those given by the core theory is, in our view, a significant advantage over the measure-theoretic framework.

4.2 Tutubalin's comments on quantum mechanics

Tutubalin explains his dissatisfaction with our discussion of quantum mechanics as follows:

> In connection with quantum mechanics, I will note that testing a physical theory is a complicated matter and goes rather deep in physics. Testing is never done directly by observing the results of an experiment: one must first put forward alternatives to the theory being tested. This problem cannot be solved by making game-theoretic bets.

We have no deep disagreement with these comments, but it is not clear to us how they bear on the points we made in our book.

We could quibble with the assertion that a theory cannot be tested without an alternative being put forward. From our point of view, any strategy for Skeptic that avoids risking bankruptcy is a test. We agree with Tutubalin, however, about how science is and should be practiced. Even if it fails certain tests, and even if it fails them repeatedly, an established physical theory will not be rejected until there is a better theory to take its place—a theory that can correct the established theory's failures while replicating its successes.

The probabilistic predictions of quantum mechanics have been thoroughly tested, and the consistency with which they have been confirmed is one important aspect of the theory's success. But as Tutubalin points out, there is much more to be said about testing a theory. Moreover, once a theory is well confirmed, its predictions interest us primarily because they tell us what to expect, not because they offer new opportunities for testing.

5 Scaling the Efficient Market Hypothesis

We conclude by discussing a question that Tutubalin raises concerning the scaling of the efficient market hypothesis.

5.1 A market game

We are now concerned with the game played by a speculator in a financial market. To fix a setting for our rather general discussion, we consider just two players: Market, who replaces both Forecaster and Reality, and Speculator, who replaces Skeptic:

Players: Market, Speculator
Protocol:
$\mathcal{K}_0 := \alpha$.
FOR day $n = 1, 2, \ldots, N$:
Market announces opening prices.
Speculator announces his purchases at these prices.
Market announces closing prices.
Speculator liquidates his holdings.
$\mathcal{K}_n := \mathcal{K}_{n-1} +$ Speculator's net gain or loss for the day.

In the context of this game, we speak of the "efficient market hypothesis" rather than of "Cournot's principle". As usual, it says that if Speculator's initial capital α is positive, then he cannot multiply it by a large factor without risking bankruptcy. We usually adopt this hypothesis for a financial market only if prices and capital are measured relative to the value of some market index such as the S&P 500. In this case, multiplying one's initial capital by a large factor without risking bankruptcy can be called "beating the market". Multiplying initial capital by a factor of ten without risking bankruptcy, for example, can be called "beating the market by a factor of ten". We will use this terminology here, without always specifying the market index that serves as our *numéraire*.

5.2 How large is too large?

When we look at an actual financial market and undertake to draw practical conclusions from our efficient market hypothesis, we must face a thorny question: just exactly how large a factor of success for Speculator do we want to rule out with our efficient market hypothesis? Do we say only that Speculator will not beat the market by a factor of 10? Or do we go further? Would it be reasonable, for example, to predict that he will not even beat it by a factor of 1.25—i.e., that he will not even do 25% better than the market?

This evidently depends on the circumstances. Several important questions must be considered, such as:

- How mature is the market? How much effort is being put into finding, exploiting, and thereby eliminating opportunities for enrichment?
- How sophisticated is the particular speculator we are considering? Does he use only relatively commonplace strategies, based on widely held information? Or is he a well-financed and well-connected hedge fund?
- How far out are we willing to stick our necks? No matter how high the limit we place on the possible success of a speculator, we may be proven wrong. How much risk of being proven wrong do we want to take? If we look at many speculators, over many time periods, follows how often are we willing to be proven wrong?

The astute reader will spot the informal probability and frequency ideas in the last question, and he may protest that our so-called "game-theoretic" framework is based on frequency ideas after all. Fair enough, but frequency is only one consideration here, and it does not enter into the formalism, which begins only after we decide how large a factor of success for Speculator we will rule out. The game-theoretic mathematics is always relative to this factor; it says that if Speculator cannot beat the market by more than the factor K, then certain events will not happen, and the larger K the fewer such events there are.

In practice, conclusions that are substantively interesting may require that we be very skeptical about Speculator's potential for success. In [3], for example, we found that the assumption that Speculator cannot beat the S&P 500 by a factor of 2 is not strong enough to yield a well-fitting game-theoretic capital asset pricing model for the average returns of stocks and simple portfolios. In order to get good results in this case, we may need an efficient market hypothesis that says that a speculator who is limited to buying or shorting stocks and simple portfolios cannot beat the S&P 500 even by a factor of 1.25 or 1.1. In any given period, a fair number of portfolios will beat the S&P 500 by more than this, and so we know that our assumption and therefore our prediction about average returns will sometimes be wrong. But they may still be right most of the time.

On p. 369 of our book, we note that a prediction based on the assumption that Speculator cannot beat the S&P 500 by a factor of 10 would have been violated by a speculator who held Microsoft stock between 1986 and 2000, and we excuse this violation with the following sentence:

> The violation is hardly surprising, because Microsoft's performance between 1986 and 2000 was in the top tenth of anyone's range of expectations.

Tutubalin expresses some puzzlement about our choice of words:

> The authors explain [the violation] by saying that Microsoft is clearly in the top 10% of all companies with respect to the growth rate of share prices, but it is not quite clear how game-theoretic probability is related to such a comparison, which probably refers to frequentist probability.

His puzzlement is justified, for "top tenth of anyone's range of expectations" hardly has a clear meaning and may suggest a stronger connection with frequency ideas than we want to advocate.

We hope, however, that the explanations we have just given make our meaning clear. Buying stock in a small firm such as Microsoft and holding it for 15 years would have been a very long shot in 1986, and we do not find it disturbing for our predictions to be violated by events so rare as the spectacular success of this particular long shot.

5.3 Upper probability and frequency

We can clarify this matter further using the notion of upper probability. The upper probability of a property E of Market's moves is

$$\overline{\mathbb{P}}\,E := \inf\{\alpha \mid \text{Speculator has a strategy that guarantees}$$
$$\mathcal{K}_N \geq 1 \text{ if Market satisfies } E \text{ and } \mathcal{K}_N \geq 0 \text{ otherwise}\}.$$

The strategies considered in this definition do not risk bankruptcy. So Market can allow E to happen only at the price of permitting Speculator to beat the market by $1/\overline{\mathbb{P}}\,E$. An upper probability of 0.1 corresponds to letting Speculator beat the market by a factor of 10, an upper probability of 0.8 corresponds to letting Speculator beat the market by a factor of 1.25, etc. We can summarize this point by saying that under the efficient market hypothesis, a small value of $\overline{\mathbb{P}}\,E$ means E is unlikely to happen. Now we see that "small" is a relative term. If we do not think Speculator can do 25% better than the market, then 0.8 is a "small" upper probability.

We can derive a one-sided connection between upper probability and frequency from the efficient market hypothesis. Consider a sequence of successive events, each with upper probability not exceeding some threshold δ. (When we say events are successive, we mean that each is settled before the upper probability of the next is assessed.) Then the one-sided strong law of large numbers (Proposition 3.4, on p. 73 of our book) implies that the proportion of the events that happen will not exceed δ if the efficient market hypothesis is not violated (cf. Proposition 15.6 on p. 365). So we may say that successive events with upper probability of 0.1 should happen no more than 10% of the time.

This thought lay behind our vague and evidently misleading explanation of why we should not be surprised that Microsoft's performance happened in spite of having upper probability of only 10%. The point is that on general principles we can claim that successive events with upper probability 10% should not happen more than 10% of the time; we have no general argument that implies they will happen less often than this. So when an event with upper probability 10% happens we should be no more surprised, in the absence of any more specific information about its improbability, than we would be by the happening of an event that happens 10% of the time in repeated trials. (We did not compare Microsoft with other stocks in the same time period, as Tutubalin's words might suggest. The simultaneous performance of different stocks does not constitute successive events.)

In a specific context, however, we may feel that certain events will be much rarer than indicated by their upper probability. We have already given one example: in very efficient markets, we might expect that it will be rare for a stock or portfolio to beat the market by more than 25%, even though this event has an upper probability of 0.8.

References

[1] Roland Omnes. *Quantum Philosophy: Understanding and Interpreting Contemporary Science*. Princeton University Press, Princeton, 1999.

[2] Glenn Shafer. Additional comments on Tutubalin's review. In preparation. Game-Theoretic Probability Project Working Paper No. 3. To appear at `http://www.cs.rhul.ac.uk/~vovk/book/`, 2002.

[3] Vladimir Vovk and Glenn Shafer. The game-theoretic capital asset pricing model. Game-Theoretic Probability Project Working Paper No. 1. Can be downloaded from `http://www.cs.rhul.ac.uk/~vovk/book/`, 2001.