

January 16, 2016

Munich Center for Mathematical Philosophy
Department of Statistics
Ludwig-Maximilians-University Munich
15–21 March 2016

Ten Lectures on the Philosophy of Game-Theoretic Probability

Glenn Shafer, Rutgers Business School, www.glennshafer.com

The game-theoretic foundation for probability, launched by a 2001 book by Vladimir Vovk and myself [24], broadens the established measure-theoretic foundation for probability. These lectures will provide an elementary introduction to the game-theoretic foundation, with an emphasis on its philosophical implications and its applications to the assessment of evidence. I hope that they will stimulate others to explore the philosophical questions the new foundation raises and the new applications it suggests. I also hope for feedback about how to relate the new foundation to the ways various disciplines, including philosophy, now think about probability.

The measure-theoretic foundation was put in definitive form by Andrei N. Kolmogorov in 1933 [9, 32] and made to work in continuous time by Joseph L. Doob in 1953 [8]. It has enabled probability theory to flourish as a mathematical enterprise, because it gives purely mathematical form to the rich language of probability, allowing mathematicians to use this language while leaving aside philosophical issues that arise whenever probability is used outside mathematics. Yet the boundary that Kolmogorov and Doob drew around their mathematics may have been too tight. While fencing out philosophy, they also excluded aspects of probability that are important in applications and can be mathematized.

Variables within Kolmogorov's framework all have probability distributions; they are *random variables*. Variables that do not have probability distributions remain outside the framework; they are *not random*. So most work that uses probability, even if it remains theoretical, lies only partly within the framework. In mathematical statistics, some dependent variables may be random only conditional on the values of independent variables that cannot be considered random for a variety of reasons. In control theory and Markov decision theory, some variables may be random only conditional on a policy for making control decisions. In quantum mechanics, the randomness (and even existence) of some measurements depends on what an observer decides to observe. Statisticians, engineers, physicists, and other scientists have found various ways to bring

mathematical results legitimized within Kolmogorov's framework to bear on the settings where such non-random variables also play a role, but the framework itself does not encompass these researchers' mathematics.

Modern game theory was not available to Kolmogorov in 1933; it emerged only with von Neumann and Morgenstern's classical book in 1944 [38]. But it is now recognized as a legitimate branch of pure mathematics, fully as rigorous even if less prestigious than older and more austere branches such as functional analysis. Subsequent to my 2001 book with Vovk, a variety of authors have demonstrated that the game-theoretic generalization of Kolmogorov's framework can retain its accomplishments while broadening their application; see especially the working papers at www.probabilityandfinance.com.

Many of the papers posted at www.probabilityandfinance.com, including some of those listed below as references for particular lectures, are highly mathematical. But the lectures will assume only an elementary understanding of probability theory.

Lecture 1. The game-theoretic definition of probability.

Tuesday, March 15, 11:30–13:00

In 1654, Pascal and Fermat gave competing solutions to the problem of dividing the stakes when a game is interrupted. Fermat's combinatorial solution presaged modern measure-theoretic probability. Pascal's solution by backward induction presaged modern game-theoretic probability.

Modern game-theoretic probability involves a game with three players. Forecaster offers bets to Skeptic, Skeptic selects from the offers, and then Reality decides the outcome. They play a finite or infinite number of rounds. An *event* is a property of the moves by Forecaster and Reality. The *probability* of an event is the capital Skeptic must risk in order to win a unit of capital if the event happens. These simple definitions suffice as a foundation for classical probability theory (Markov's inequality, the law of large numbers, etc.).

When the offers made by Forecaster fall short of defining full probability measures, we get only upper probabilities and upper expectations. Natural axioms for the offers lead to axioms for upper expectations, and if the axioms hold for events on each round, they also hold for events in the whole course of play.

The axioms for upper expectations in game-theoretic probability are more or less equivalent to axioms used by Peter Walley's theory of imprecise probabilities. Walley did not, however, formulate a game; his logic was concerned only with attitudes of a single individual, who resembles game-theoretic probability's Forecaster. Game-theoretic probability considers three players, and in particular it has at its disposal the powerful notion of a strategy for Skeptic.

References: The basic ideas of game-theoretic probability are explicated in [28] and [31]. Walley's basic ideas were presented in [44], but this book is out of print and not widely available; more up-to-date and accessible references include [1] and [36].

Lecture 2. Frequentism.

Tuesday, March 15, 14:00–15:30

The founders of the measure-theoretic foundation for mathematical probability, including Emile Borel, Paul Lévy, Maurice Fréchet, Andrei Kolmogorov, and Joseph Doob, identified probability with long-run frequency. But the link between measure theory and frequency has proven philosophically troublesome. Many scholars believe that a story about propensities or multiple worlds is needed in order to relate frequencies to measure theory.

The game-theoretic foundation connects probability with observed frequencies in a way more faithful to the thinking of Borel and Kolmogorov and the older traditions they represented. The connection is made by Skeptic's strategies for betting against hypothesized upper expectations. Simple strategies for Skeptic will multiply the capital they risk by a large factor if observed frequencies do not conform to these upper expectations.

Borel, Kolmogorov, and their contemporaries used Cournot's principle to connect probability with frequency. Cournot's principle says that an event of small probability will not happen. The game-theoretic version says that Skeptic will not multiply the capital he risks by a large factor.

In this lecture, I review how classical probabilists such as Borel and Kolmogorov and classical statisticians such as Jerzy Neyman explained their frequentist philosophy of probability and argue that the game-theoretic understanding of statistical testing is a natural elaboration of their thinking. The game-theoretic understanding is especially successful in helping us navigate the generalization from the traditional picture of repeated trials to the modern picture of stochastic processes.

References: My review of the older authors will follow [26] and [32]. My discussion of Neyman's views will rely mainly on [13] and [14].

Lecture 3. Dynamic martingale testing.

Wednesday, March 16, 11:30–13:00

A *martingale* is the capital process determined by a strategy for Skeptic—it tells what Skeptic's capital will be at the end of each round as a function of the previous moves by Skeptic's opponents. A martingale is nonnegative if and only if the strategy risks no more than its initial capital. So game-theoretic testing means testing using a nonnegative martingale. We reject proposed probabilities (or upper expectations) when the nonnegative martingale we are using to test them becomes very large. Because the value of a nonnegative martingale can come down after going up, this is a dynamic concept of testing.

Jean Ville invented martingales in the 1930s to show the inadequacy of the frequentist theory created by Richard von Mises and made algorithmic by Abraham Wald. Starting in the 1960s, Ville's idea was further developed in the algorithmic direction by Per Martin-Löf, Claus-Peter Schnorr, and A. Philip Dawid. Martingales were brought into measure-theoretic probability by Doob

in the 1940s and subsequently became important in mathematical statistics. Mathematical statistics has been slow, however, to take up dynamic testing.

The static approach to testing begins with a test statistic; a hypothesis is rejected if the statistic is too much larger than expected under the hypothesis. This focuses our attention on the probability, under the hypothesis, that the test statistic would have equaled or exceeded the value observed; this is now called the *p-value*. The use of p-values has been criticized by classical statisticians, who insist that a rejection level (the *significance level*) be fixed in advance, and by Bayesians, who find Bayes factors, which are usually more conversative, more convincing.

When we test using a nonnegative martingale that begins with unit capital, the evidence against the hypothesis is measured simply by the martingale's current value, or equivalently by its inverse: by Markov's inequality, a value of 1000 or more for the martingale has (upper) probability of 1/1000 or less. We may be tempted to cheat by using instead the greatest value attained by the martingale so far. This cheating is homologous, mathematically and philosophically, to the cheating involved in using a p-value instead of a significance level fixed in advance. (We cheat by pretending that we had planned in advance to stop just where the highest value of the martingale occurred, just as we cheat by pretending that we fixed in advance a significance level equal to the p-value actually observed.)

As we will see, there are game-theoretic ways to discount the value of the maximum of a martingale so far (or a p-value) so that its evidential force is not lost entirely but the cheating is taken into account.

References: Dynamic game-theoretic testing is explained in [7], and testing using nonnegative martingales in measure-theoretic probability is explained in [30]. For details about the role of martingales in algorithmic randomness from von Mises to Schnorr, see [3, 37, 43]. For a historical review of the paradoxical behavior of martingales when they are not required to be nonnegative (or at least bounded below), see [4]. For the history of the term *p-value*, see John Aldrich's comments at <http://jeff560.tripod.com/p.html>.

Lecture 4. How to forecast.

Wednesday, March 16, 14:00–15:30

So far we have put ourselves in the shoes of Skeptic, who establishes the empirical meaning of probabilities (and upper expectations) by testing them. Now let us put ourselves in the shoes of Forecaster. Can we make probability forecasts that pass Skeptic's tests?

As it turns out, Forecaster can defeat any particular strategy for Skeptic, provided only that each move prescribed by the strategy varies continuously with respect to Forecaster's previous move. Forecaster wants to defeat more than a single strategy for Skeptic; he wants to defeat simultaneously all the strategies Skeptic might use. But as we will see, Forecaster can often amalgamate the strategies he needs to defeat by averaging them, and then he can play against the average. This is called *defensive forecasting*. Defeating the average may

be good enough, because when any one of the strategies rejects Forecaster's validity, the average will reject as well, albeit less strongly.

Probabilities can be estimated consistently using a random sample. But this well known fact assumes that observations are independently drawn from a fixed probability distribution. Defensive forecasting, in contrast, gives probabilities that pass statistical tests without using any advance knowledge about how reality will behave.

The thesis that adaptive methods of forecasting can often do a good job without advance scientific knowledge is not new, but this theoretical validation of the thesis has fundamental implications for the meaning of probability. It reveals that the crucial step in placing an evidential question in a probabilistic framework is its placement in a sequence of questions. Once we have chosen the sequence, good sequential probabilities can be given, and the validation of these probabilities by experience signifies less than commonly thought.

This insight lends weight to Jerzy Neyman's proposal to found classical statistics on the concept of *inductive behavior*. It also clarifies the role of theories of evidence, such as the Dempster-Shafer theory and even radical subjectivist Bayesianism. These theories have their place in situations where there is no agreement on a sequence of comparable questions.

References: Defensive forecasting is explained in detail in [27] and [42]. For additional references to related work, see page 95 of [5]. Neyman explained his concept of inductive behavior in [12]; see also [11].

Lecture 5. Subjective probability.

Thursday, March 17, 11:30–13:00

The game-theoretic picture permits a variety of subjective interpretations of its probabilities or its upper probabilities and upper expectations.

In the interpretation famously advocated by Bruno de Finetti and generalized by Peter Walley, having a subjective probability means being willing to bet at the odds it implies; or at least considering such bets acceptable or desirable. In our game-theoretic picture, we might imagine the Forecaster taking this attitude towards the bets he offers, though the logic of testing does not require him to do so.

Another way of giving a subjective interpretation to the upper expectations in the game-theoretic picture, which does rely on the logic of testing, is to assert that a person buying payoffs at these prices will not be able to refute them by multiplying the stake he risks by a large factor. This has many variations, some of which may sound more subjective than others. I might assert a belief that no betting strategy I can devise with the information I have will refute given probabilities (or probabilities to be obtained from a theory or a Forecaster whom I respect). Or I might assert that no one, regardless of their information, can do so.

These various subjective interpretations lead to various responses to the historically vexed question of how subjective probabilities should change as

information is obtained. If we are in Forecaster's shoes, within the perfect-information game, then the question is how Forecaster produces his forecasts. If we are standing outside the picture and do not fully observe it, then we must make other judgements. We may or may not make judgements that conform to Laplace's *principle of inverse probability*, according to which the probabilities we attribute to different hypotheses after observing an event should be proportional to the probabilities the hypotheses attribute to the event. We may or may not make judgements that conform to Thomas Bayes's second argument for conditional probability.

References: In his *History of Statistics*, Stephen M. Stigler argues that Laplace initially arrived at his principle through a mathematical error (pages 113–117 of [34]). My 1982 critique of Bayes's second argument for conditional probability is [19] and was discussed by Andrew Dale in [6]. My arguments for the constructive nature of subjective probability were published in the 1980s in [17, 18, 21].

Lecture 6. Dempster's rule of combination.

Thursday, March 17, 14:00–15:30

As explained in my 1976 book, theory of belief functions (also known as the Dempster-Shafer theory) is a generalization of probability theory; a belief function is a set function more general than a probability measure but whose values can still be interpreted as degrees of belief. Dempster's rule of combination is a rule for combining two or more belief functions; when the belief functions combined are based on independent bodies of evidence, the rule corresponds intuitively to the pooling of evidence. As a special case, the rule yields a rule of conditioning that generalizes the Bayesian rule for updating probability measures.

But how do we know that two bodies of evidence are *independent*? Because the independence of events or variables has a mathematical definition within measure-theoretic probability, many critics have demanded some analogous definition for bodies of evidence. But in the theory of belief functions, independence is a constructive judgement, not a something that is already mathematically true. When we say that two bodies of evidence from which we have constructed belief functions bearing on the same question are independent, we are saying that the probabilistic judgements made from the one body of evidence are not changed by the other body of evidence, except insofar as the two conflict directly. As I have argued, a similar constructive judgement underlies Bayesian updating: in order to change our probability for B after learning that A is true from $P(B)$ to $P(A\&B)/P(A)$, we must make a judgement that learning A does not tell us anything else about B .

The game-theoretic foundation for probability and its logic of testing allow us to explain more clearly the content of these constructive judgments. In both the Bayesian and belief-function cases, we are judging that the new information does not change our belief that we cannot devise betting strategies that will defeat certain probabilistic predictions. We do not believe that an opponent

(Skeptic) can multiply the capital he risks by a large factor betting at the odds they offer.

In the Bayesian case, when we say that we have learned A and nothing more, the *nothing more* means nothing more that would help us defeat the odds offered. We can show that for any strategy for Skeptic that uses initial probabilities $P(A)$ and $P(A\&B)$ and then $P(A\&B)/P(A)$ after A is known, there exists a strategy with the same cost and payoff that uses only the initial probabilities $P(A)$ and $P(A\&B)$. So *provided that nothing else becomes known that helps Skeptic choose his strategy*, his inability to beat the initial probabilities implies an inability also to beat the new conditional probability. The caveat that nothing else becomes known that helps him bet is essential, and this must be supplied by exogenous judgement—it is not supplied by the mere definition of conditional probability.

In the case of Dempster-Shafer theory, using a multivalued mapping that moves from game-theoretic probabilities (or upper probabilities or upper expectations) to Dempster-Shafer degrees of belief requires a similar judgement of irrelevance. Learning the multivalued mapping does not, we believe, help an opponent defeat the probabilities. Combining two belief functions by Dempster's rule then requires further judgments of irrelevance: learning the two multivalued mappings does nothing more to help the opponent. Again, these are exogenous judgements.

The same logic justifies the intuition that originally motivated Laplace's principle of inverse probability. Our observation y of a quantity θ has an error e : $y = \theta + e$. We begin with beliefs about (probabilities for) e , and having observed y , we continue to have the same beliefs about $y - \theta$. This "continuing to believe" is justified by a game-theoretic judgment of irrelevance, the judgement that our learning y does not help an opponent defeat our probabilities. The counterexamples that we now have to this judgment (on the margin, learning y does sometimes help the opponent) prevent us from making it too freely but do not rule out ever making it.

References: Dempster's rule is explained, of course, in my 1976 book [16]. The game-theoretic interpretation is explained in [25] and [29].

Lecture 7. Constructive decision theory.

Friday, March 18, 11:30–13:00

In the 1980s, in work only now being published, I argued that decision making can be approached in a constructive way that meshes with the constructive interpretation of Dempster-Shafer theory.

From a constructive point of view, Bayesian decision theory breaks a decision problem down by decoupling our consideration of utilities for different outcomes from our consideration of probabilities. We first assign a utility to each possible outcome. For each action we might take, we consider the probabilities it will produce the different outcomes and combine these with the utilities to obtain an expected utility for the action. But as Leonard J. Savage pointed out in 1954, in his celebrated *Foundation of Statistics*, the decoupling is illusory. No matter how much detail we put into the description of an outcome, our happiness with

it will still depend on additional contingencies, and hence the utility we assign it must itself be interpreted as an expected utility. Savage called this *the problem of small worlds*.

Taking a resolutely constructive but less utilitarian view, I suggested that some decision problems can be better analyzed by constructing goals. When we think of a goal as something we invent, it is natural to put a weight on a goal without supposing that this weight is an expected value of future satisfaction or an average over more detailed preferences about the future state of the world. It is then also natural to score an action in terms of the degree to which it assures the goal is achieved, in the spirit of Dempster-Shafer theory. This approach is not necessarily game-theoretic, but it becomes almost necessary in the framework of game-theoretic framework, where there are no exogenous probabilities.

In 2007, David Krantz and Howard Kunreuther proposed a very similar approach to decision making and used it to rationalize the way people pursue insurance-related goals.

References: My unpublished work on constructive decision theory is [20]; my related critique of Savage's understanding of small worlds [15] appeared in 1986 [22]. Krantz and Kunreuther's work is at [10].

Lecture 8. Causality as regularity.

Friday, March 18, 14:00–15:30

Most of the literature on probabilistic causality uses *variables* as the unit of analysis. Certain variables are seen as causes of other variables and events. Within game-theoretic probability, this seems wrong-headed, because variables and events are usually global objects—summaries, as it were, of the entire course of play. Formally, an *event* is a property of the whole sequence of moves by Forecaster and Reality, and a *variable* is a real-valued function of this sequence. It is more natural to talk about the moves by Reality as causes contributing to how these global events and variables come out.

My 1996 book on causality did not use the language of games, but it took a probability tree—or game tree, as I might now say—as basic and analyzed causal relations among variables as properties of this tree—regularities in how Reality's later moves are related to her earlier moves. This analysis can be translated into the language of game-theoretic probability.

This is a further illustration of the diversity of ways in which game-theoretic probability can be used. We have seen examples of probability games in which we play the role of Skeptic or Forecaster, and in both cases we see all the moves by all the players. But when we are studying causality, we are outside the relevant game. The players may see each others' moves, but we do not see Reality's moves, which are the causes we seek to learn about.

References: My 1996 book on probabilistic causality was entitled *The Art of Causal Conjecture* [23]. Chapter 1 and some related papers are available at <http://www.glennshafer.com/books/acc.html>.

Lecture 9. Game-theoretic finance I.

Monday, March 21, 11:30–13:00

The efficient-market hypothesis, as formulated by Eugene Fama in the 1960s and repeated in introductions to finance theory today, is admittedly circular; this is called the *joint hypothesis problem*. Fama’s hypothesis comes down to saying that today’s price is the expected value of tomorrow’s price, but if arbitrage is not possible then this true for some probability distribution, and the hypothesis cannot be tested until the probability distribution is specified.

The game-theoretic version of Cournot’s principle provides a statement of the efficient-market hypothesis that avoids this circularity. To see this, we have the market play the role of both Forecaster and Reality; today’s price being the prediction of tomorrow’s price—the price at which we can buy or sell it. We imagine two players who buy and sell at these prices: Investor and Skeptic, who observes Investor’s moves and buys and sells to form bets on how Investor will perform. The efficient-market hypothesis is that Skeptic will not multiply the capital he risks at a large factor.

From this efficient-market hypothesis we obtain, without any “stochastic” hypotheses, versions of many of the standard results of finance theory, including the capital asset pricing model (CAPM), the equity premium, and the \sqrt{t} effect (variation resembling Brownian motion).

References: An elementary explanation of the \sqrt{t} effect is given in [40]. CAPM is derived in [39], and the lead-lag effects are discussed in [45]. See also the Tokyo papers at <http://www.probabilityandfinance.com/articles/index.html> and the recent book by Elie Ayache [2].

Lecture 10. Game-theoretic finance II.

Monday, March 21, 14:00–15:30

In our 2001 book on game-theoretic probability and finance, Vovk and I used non-standard analysis to study probability-free finance in continuous time. A more successful approach was proposed in 2007 by Kei Takeuchi, Masayuki Kumon, and Akimichi Takemura. In their formulation, the market moves continuously, but an investor trades in discrete times, which can depend on the past path of the market. The investor’s strategy can be described by saying that he divides his capital into countably many accounts, which trade with higher and higher frequency.

References: The work by Takeuchi et al. is at [35]. The emergence of Brownian motion from this picture is elegantly explained in [41].

References

- [1] Thomas Augustin, Frank P. A. Coolen, Gert de Cooman, and Matthias C. M. Troffaes, editors. *Introduction to Imprecise Probabilities*. Wiley, Chichester, 2014.

- [2] Elie Ayache. *The Medium of Contingency: An Inverse View of the Market*. Wiley, 2015.
- [3] Laurent Bienvenu, Glenn Shafer, and Alexander Shen. On the history of martingales in the study of randomness. *Electronic Journal for History of Probability and Statistics*, 5(1), 2009. www.jehps.net
- [4] Bernard Bru, Marie-France Bru, and Kai Lai Chung. Borel and the St. Petersburg martingale. *Electronic Journal for History of Probability and Statistics*, 5(1), 2009. www.jehps.net
- [5] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge, 2006.
- [6] Andrew I. Dale. *A History of Inverse Probability from Thomas Bayes to Karl Pearson*. Springer, 1999.
- [7] A. Philip Dawid, Steven de Rooij, Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Insuring against loss of evidence in game-theoretic probability. Working Paper #34 at <http://www.probabilityandfinance.com/articles/34.pdf>, 2015.
- [8] Joseph L. Doob. *Stochastic Processes*. Wiley, New York, 1953.
- [9] Andrei N. Kolmogorov. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. In *Ergebnisse der Mathematik*, volume 2, pages 195–262. Springer (for *Zentralblatt für Mathematik*), Berlin, 1933. An English translation by Nathan Morrison appeared under the title *Foundations of the Theory of Probability* (Chelsea, New York) in 1950, with a second edition in 1956.
- [10] David H. Krantz and Howard C. Kunreuther. Goals and plans in decision making. *Judgment and Decision Making*, 2(3):137–168, 2007. <http://journal.sjdm.org/jdm7303b.pdf>
- [11] Erich L. Lehmann. Neyman’s statistical philosophy. *Probability and Mathematical Statistics*, 15:29–36, 1995. <http://www.math.uni.wroc.pl/~pms/publications.php?nr=15>
- [12] Jerzy Neyman. , ‘Inductive behavior’ as a basic concept of philosophy of science. *Review of the International Statistical Institute*. 25:7–22, 1957.
- [13] Jerzy Neyman. Indeterminism in science and new demands on statisticians. *Journal of the American Statistical Association*, 55:625–639, 1960.
- [14] Jerzy Neyman. Frequentist probability and frequentist statistics. *Synthese*, 36:97–131, 1977.
- [15] Leonard J. Savage. *The Foundations of Statistics*. Wiley, 1954. Second edition, Dover, 1972.
- [16] Glenn Shafer. *A Mathematical Theory of Evidence*. Princeton, 1976.

- [17] Glenn Shafer. Constructive probability. *Synthese*, 48:1–60, 1981. http://www.glennshafer.com/assets/downloads/articles/article05_constructive.pdf.
- [18] Glenn Shafer. Two theories of probability, PSA 1978, Vol. 2, pp. 441–464. Peter D. Asquith and Ian Hacking, eds., 1981. http://www.glennshafer.com/assets/downloads/articles/article07_TwoTheories1981.pdf
- [19] Glenn Shafer. Bayes’s two arguments for the rule of conditioning. *Annals of Statistics*, 10:1075–1089, 1982. <http://projecteuclid.org/euclid.aos/1176345965>
- [20] Glenn Shafer. Constructive decision theory. To appear in *International Journal of Approximate Reasoning*. <http://www.glennshafer.com/assets/downloads/decision.pdf>
- [21] Glenn Shafer and Amos Tversky. Languages and designs for probability judgement. *Cognitive Science*, 9:309–339, 1985. http://www.glennshafer.com/assets/downloads/articles/article19_languages.pdf
- [22] Glenn Shafer. Savage revisited (with discussion). *Statistical Science* 1:463–501, 1986. <http://projecteuclid.org/euclid.ss/1177013504>
- [23] Glenn Shafer *The Art of Causal Conjecture*. MIT, 1996.
- [24] Glenn Shafer and Vladimir Vovk. *Probability and Finance: It’s Only a Game!* Wiley, New York, 2001.
- [25] Glenn Shafer, Peter R. Gillett and Richard B. Scherl. A new understanding of subjective probability and its generalization to lower and upper prevision. Working Paper #3 at <http://www.probabilityandfinance.com/articles/03.pdf>, 2002.
- [26] Glenn Shafer. From Cournot’s principle to market efficiency. Working Paper #15 at <http://www.probabilityandfinance.com/articles/15.pdf>, 2006.
- [27] Glenn Shafer. Game-theoretic probability and its uses, especially defensive forecasting. Working Paper #22 at <http://www.probabilityandfinance.com/articles/22.pdf>, 2007.
- [28] Glenn Shafer, Vladimir Vovk, and Roman Chyčhyla. How to base probability theory on perfect-information games. Working Paper #32 at <http://www.probabilityandfinance.com/articles/32.pdf>, 2009.
- [29] Glenn Shafer. A betting interpretation for Dempster-Shafer degrees of belief. Working Paper #31 at <http://arxiv.org/abs/1001.1653>, 2010.

- [30] Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Test martingales, Bayes factors, and p-values. Working Paper #33 at <http://www.probabilityandfinance.com/articles/33.pdf>, 2014.
- [31] Glenn Shafer, Vladimir Vovk, and Akimichi Takemura. Lévy’s zero-one law in game-theoretic probability. Working Paper #29 at <http://arxiv.org/abs/0905.0254>, 2010.
- [32] Glenn Shafer and Vladimir Vovk. The origins and legacy of Kolmogorov’s *Grundbegriffe*. Working Paper #4 at <http://www.probabilityandfinance.com/articles/04.pdf>, 2013.
- [33] Glenn Shafer. When to call a variable random. Working Paper #41 at <http://www.probabilityandfinance.com/articles/41.pdf>, 2015.
- [34] Stephen M. Stigler. *The History of Statistics: The Measurement of Uncertainty before 1900*. Harvard, 1986.
- [35] Kei Takeuchi, Masayuki Kumon, and Akimichi Takemura. A new formulation of asset trading games in continuous time with essential forcing of variation exponent. *Bernoulli*, 15:1243–1258, 2009. <http://arxiv.org/abs/0708.0275>
- [36] Matthias C. M. Troffaes and Gert de Cooman. *Lower Previsions*. Wiley, 2014.
- [37] Jean Ville. *Étude critique de la notion de collectif*. Gauthier-Villars, Paris, 1939.
- [38] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton, 1944, 1947, 1953.
- [39] Vladimir Vovk and Glenn Shafer. The game-theoretic Capital Asset Pricing Model. Working Paper #1 at <http://www.probabilityandfinance.com/articles/01.pdf>, 2002.
- [40] Vladimir Vovk and Glenn Shafer. A game-theoretic explanation of the \sqrt{dt} effect. Working Paper #5 at <http://www.probabilityandfinance.com/articles/05.pdf>, 2003.
- [41] Vladimir Vovk. Continuous-time trading and the emergence of probability. Working Paper #28 at <http://www.probabilityandfinance.com/articles/28.pdf>, 2011.
- [42] Vladimir Vovk, Akimichi Takemura, and Glenn Shafer. Defensive forecasting. Working Paper #8 at <http://www.probabilityandfinance.com/articles/08.pdf>, 2005.
- [43] Abraham Wald. Die Widerspruchfreiheit des Kollektivbegriffes der Wahrscheinlichkeitsrechnung. *Ergebnisse eines Mathematischen Kolloquiums*, 8:38–72, 1937.

- [44] Peter Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, 1991.
- [45] Wei Wu and Glenn Shafer. Testing lead-lag effects under game-theoretic efficient market hypotheses. Working Paper #23 at <http://www.probabilityandfinance.com/articles/23.pdf>, 2007.