


The diachronic Bayesian

Vladimir Vovk



The Game-Theoretic Probability and Finance Project

Working Paper #64

First posted August 24, 2023. Last revised August 24, 2025.

Project web site:

<http://www.probabilityandfinance.com>

Abstract

It is well known that a Bayesian probability forecast for all future observations should be a probability measure in order to satisfy a natural condition of coherence. The main topics of this paper are the evolution of the Bayesian probability measure and ways of testing its adequacy as it evolves over time. The process of testing evolving Bayesian beliefs is modelled in terms of betting. The resulting picture is adapted to forecasting several steps ahead and making almost optimal decisions.

This paper has also been published as an [arXiv report](#).

Contents

1	Introduction	1
1.1	Completely predictive approach	2
1.2	Notation and terminology	2
1.3	Dramatis personae	3
2	Basic prediction picture	4
3	Testing probability forecasts	5
3.1	Comparison with Bayesian conditioning	7
3.2	Merging Sceptic's opponents	8
4	Measure-theoretic picture	9
4.1	Validity of the game-theoretic picture	9
4.2	Validity of the measure-theoretic picture	10
5	Predicting K steps ahead	11
6	Bayesian decision making	12
6.1	An optimality result for the Bayes decision strategy	12
6.2	Predicting $K < N$ steps ahead is essential for our statement of optimality	16
7	Conclusion	18
	References	18

A	Proofs	21
A.1	Proof of Proposition 3.2	21
A.2	Proof of Theorem 4.2 and Proposition 4.1	21
A.3	Game-theoretic probability	23
A.4	Proof of Theorem 6.5	24
A.5	Proof of Proposition 6.8	29
A.6	Proof of Proposition 6.9	29
B	Protocols in terms of ideal futures	30
B.1	Ideal futures contracts	31
B.2	General testing protocol	31
B.3	Simplification	33
C	Mechanics of futures trading	34
D	Radical probabilism	36
D.1	Additive picture	37
D.2	Multiplicative picture	37
D.3	Radical probabilism and reality	38
E	Internal and external coherence	39

1 Introduction

Consider a Bayesian forecaster predicting future observations. Two standard examples, which can be considered as the opposite ends of a spectrum, are where the observations are outcomes of coin tosses (for a possibly biased coin) and where the observations are “dry” or “rain” for a number of consecutive days. Let us take the standard Bayesian position, due to [de Finetti \[1937, 2017\]](#) and discussed by, e.g., [Bernardo and Smith \[2000, Sect. 4.1\]](#), that the Bayesian’s beliefs about the future observations should be encoded as a probability measure on the sequences of observations.

Therefore, we assume that at each point in time the Bayesian has a probability measure representing his beliefs for the future observations. In this paper we are interested in how the Bayesian’s probability measure changes over time. We will formalize such a prediction picture as various prediction protocols. Some prediction protocols are comprehensive in that they record **all** information that the Bayesian gets, and then it is natural for the Bayesian to condition on the new information in the usual sense of probability theory [[Kolmogorov, 1933](#), Sect. I.4]. This procedure for updating the Bayesian’s beliefs is known as “Bayesian conditioning” [[Bayes, 1763](#), [Shafer, 1982, 2022](#)]. [Lewis \[1999\]](#) derives Bayesian conditioning as updating rule via his diachronic Dutch book argument, which relies on the prediction protocol being comprehensive. In [Sect. 2](#) we will discuss the narrowness of comprehensive prediction protocols and, therefore, of Bayesian conditioning. While devising comprehensive prediction protocols is feasible in coin-type situations, it is not in weather-type ones (cf. the first paragraph of this section).

[Section 3](#) proposes a testing protocol based on betting for the evolving Bayesian probability measure. This protocol is given in terms of observables and does not depend on Bayesian conditioning. It is very much in the spirit of the work on game-theoretic probability [[Shafer and Vovk, 2019](#), [Dawid and Vovk, 1999](#), [Vovk, 1993](#)] and the recent discussion papers by [Shafer \[2021\]](#), [Waudby-Smith and Ramdas \[2024\]](#), and [Grünwald et al. \[2024\]](#) promoting game-theoretic statistics. While standard game-theoretic probability only deals with one-step ahead prediction, this paper extends it to predicting the full future. This leads to various conceptual and technical complications.

A discussion of connections with the standard measure-theoretic picture of probability and statistics follows in [Sect. 4](#). The measure-theoretic picture will typically be an imaginary picture that does involve Bayesian conditioning (which may be happening deeply inside the imaginary picture, far from what we can observe). The main finding of this section is the equivalence of the game-theoretic and measure-theoretic pictures for finite probability spaces. However, this theoretical vindication of Bayesian conditioning is very remote from what we can observe in our prediction protocols.

[Section 5](#) adapts the testing protocol of [Sect. 3](#) to predicting K steps ahead, which generalizes the case $K = 1$ considered earlier in game-theoretic probability (e.g., [Shafer and Vovk, 2019](#), [Dawid and Vovk, 1999](#), [Vovk, 1993](#)). [Section 6](#) applies the K -steps-ahead testing protocol to making nearly optimal decisions,

and Sect. 7 concludes.

Appendixes A–E provide further information. The key one is Appendix A giving the proofs. Appendix B discusses a seemingly more general (but in fact equivalent) testing procedure, and Appendix D discusses Jeffrey’s radical probabilism in the context of testing diachronic Bayesian predictions.

This paper was inspired by Philip Dawid’s brief discussion of one-step-ahead prediction [Vovk and Shafer, 2025, Sect. 7], and its title is adapted from Dawid [1982] (being well-calibrated is an important aspect, namely the frequency aspect, of agreement with reality). Its other source is Dawid’s super-strong prequential principle [Dawid and Vovk, 1999, Sect. 5.2]. This principle requires that our testing protocol based on betting must agree with the measure-theoretic picture, regardless of the imagined data-generating distribution.

1.1 Completely predictive approach

A very common Bayesian picture is where, instead of one probability measure P over the observations, the Bayesian’s beliefs are modelled as a statistical model $\{P_\theta \mid \theta \in \Theta\}$ combined with a prior probability measure μ on Θ . A standard Bayesian point of view [Bernardo and Smith, 2000, Chap. 4] is that we should start from one probability measure P and only then, if this is more convenient, e.g., mathematically, represent it as integral $P = \int P_\theta \mu(d\theta)$. An example is the application of de Finetti’s theorem to coin tossing, guaranteeing that any exchangeable probability measure P can be represented as a mixture $\int P_\theta \mu(d\theta)$ of probability measures P_θ corresponding to independent and identically distributed observations. The case of one probability measure P is simpler and more fundamental. In Lindley’s words, “We should be concentrating not on Greek letters but on the Roman letters” (i.e., not on θ s, parameter values, but on the x s and y s, observables) [Vovk and Shafer, 2025, Sect. 7]. This view is sometimes called the completely predictive approach [Piccinato, 1986].

1.2 Notation and terminology

If a and b are finite sequences (of some elements), we write $a \subseteq b$ to mean that a is a prefix of b , and we write $a \subset b$ to mean that $a \subseteq b$ and $a \neq b$. If $a \subseteq b$, $b \setminus a$ is the sequence obtained from b by crossing out its prefix a . The concatenation of a and b is written simply as ab ; we use the same notation when a or b (or both) are elements; if B is a set of elements or finite sequences, aB stands for $\{ab \mid b \in B\}$. The length of a finite sequence a is denoted by $|a|$; in particular, $|\square| = 0$ for the empty sequence \square .

If a and b are numbers, $a \wedge b := \min(a, b)$.

In our terminology we will mainly follow Shiryayev [2016, 2019]. A *finite probability space* is a pair (Ω, P) , where Ω is a finite set, implicitly equipped with the σ -algebra \mathcal{F} of all subsets of Ω , and P is a probability measure on (Ω, \mathcal{F}) . Let us say that (Ω, P) is *positive* if the probability of each sample point is positive, $P(\{\omega\}) > 0$ for all $\omega \in \Omega$. We let \mathbb{E}_P stand for the expected value under P .

We will also use the following notation:

- $\mathfrak{P}(A)$ is the set of all positive probability measures on a finite set A ;
- if $P \in \mathfrak{P}(\mathbf{Y}^K)$ and $x \in \mathbf{Y}^k$ for a finite set \mathbf{Y} and $k \leq K$,

$$P(x) := P(x\mathbf{Y}^{K-k}); \quad (1)$$

- if $P \in \mathfrak{P}(\mathbf{Y}^K)$, $x \in \mathbf{Y}^k$, and $x' \in \mathbf{Y}^{k'}$ for a finite set \mathbf{Y} and $k + k' \leq K$,

$$P(x' | x) := \frac{P(xx')}{P(x)}$$

(under our definitions the denominator will always be positive).

A *filtration* (\mathcal{F}_n) in a finite probability space (Ω, P) , where n ranges over a contiguous set of integers, is an increasing sequence of σ -algebras in Ω , $\mathcal{F}_{n_1} \subseteq \mathcal{F}_{n_2}$ when $n_1 \leq n_2$. We say that a sequence (Y_n) of random variables in (Ω, P) is *adapted* if Y_n is \mathcal{F}_n -measurable for all n ; it is *predictable* if Y_n is \mathcal{F}_{n-1} -measurable for all n . We will mainly concentrate on finite probability spaces.

In this paper I will ignore any distinctions that are sometimes made between “forecast” and “prediction” and will regard these words as synonymous. I will never use more exotic words such as “prevision” [de Finetti, 2017, Sect. 3.1.2].

Following Shafer and Vovk [2001, 2019] I use “game-theoretic” to refer to being based on betting, and the kind of game theory involved here is the theory of perfect-information games rather than the probabilistic games studied in, e.g., economics (see, e.g., Shafer and Vovk 2019, Sect. 4.5).

1.3 Dramatis personae

These are the players in our prediction protocols (most of the protocols involve subsets of players).

- Reality (female): player who chooses sequential observations y_1, y_2, \dots , which are elements of the observation space \mathbf{Y} (assumed finite).
- Forecaster (male): player who issues probabilistic forecasts for the future observations.
- Sceptic (male): player who gambles against Forecaster’s predictions. Informally, he is trying to discredit Forecaster.
- Decision Maker (female): player who makes decisions in light of Forecaster’s predictions.

The players’ sexes are defined in Shafer and Vovk [2019]. The noun “Bayesian” will often be used as nearly synonymous with “Forecaster”, and so the Bayesian is male.

2 Basic prediction picture

We are interested in the following sequential *Bayesian prediction protocol*.

Protocol 2.1.

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

Forecaster announces $P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1})$

Reality announces the actual observation $y_n \in \mathbf{Y}$.

In this paper we only consider the case of a finite *time horizon* $N > 1$. At each step n , P_n is a prediction for the whole future $y_n y_{n+1} \dots y_N$ (sequence of length $N - n + 1$). Earlier we referred to Forecaster as “Bayesian”, in order to emphasize that his predictions are complete probability measures over the future observations (while in earlier work we often considered less complete predictions: see, e.g., [Shafer and Vovk 2019](#), Preface, point 2). Forecaster qua Bayesian is diachronic, as his beliefs evolve sequentially over time.

Protocol 2.1 does not define a game, since we have not specified the players’ goals, but we will often talk about the *plays* $P_1 y_1 \dots P_N y_N$ proceeding according to the protocol’s rules.

Let us assume, for simplicity, that the set \mathbf{Y} is finite (and equipped with the discrete σ -algebra); this will allow us to concentrate on conceptual issues avoiding technical difficulties and ambiguities (such as countable vs finite additivity). To exclude trivialities, we also assume $|\mathbf{Y}| > 1$.

In addition, we impose the requirement that $P_n(E) > 0$ unless $E = \emptyset$. This is a version of Lindley’s “Cromwell’s rule” [[Lindley, 1985](#), Sect. 6.7].

Protocol 2.1 goes beyond *Bayesian conditioning*, where we insist that, for each $n \geq 2$,

$$P_n(x) := P_{n-1}(x \mid y_{n-1}) := \frac{P_{n-1}(y_{n-1}x)}{P_{n-1}(y_{n-1})}, \quad x \in \mathbf{Y}^{N-n+1}. \quad (2)$$

Bayesian conditioning (as rule for updating beliefs) was criticised by [Hacking \[1967\]](#), since the rule ignores the cost of thinking. [Lewis \[1999\]](#) points out that “we should sometimes respond to conceptual discoveries by revising our beliefs”. However, the most straightforward reason for violating (2) is that at step n Forecaster can also learn other information apart from y_n (i.e., learn information outside the protocol). Bruno [de Finetti \[2017\]](#) stated this objection in this way:

It would be wrong... to state, or to think, in a superficial manner, without at least making sure that these explanations are implicit, that $P(E|H)$ is the probability of E once H is known. In general, by the time we learn that H has occurred, we will already have learnt of other circumstances that might also influence our judgement. In any case, the evidence that establishes that H has occurred will itself contain, explicitly or implicitly, a wealth of further detail, which will modify our final state of information and, most likely, our probabilistic judgement.

This objection was also discussed at length by [Shafer \[1985\]](#). And when presenting his diachronic Dutch book argument for Bayesian conditioning, [Lewis \[1999\]](#) emphasized that the observations on which we condition should “specify, in full detail, all the alternative courses of experience you might undergo”.

Let me give an example where Bayesian conditioning is utterly unrealistic as an updating rule, because we can’t hope to have a comprehensive protocol including all the information a real-life Bayesian has access to. Consider the standard case [\[Dawid, 1982, 2006\]](#) of a weather forecaster who issues a probability for the rain on sequentially numbered days. The observations are the actual outcomes, say 0 or 1 (encoding a dry or rainy day). In the morning of day 1 the forecaster announces a joint probability for the future observations (for days 1, 2, ...) as his forecast, and in the morning of day 2 he announces a new forecast, for days 2, 3, We can’t assume that the 0/1 observation on day 1 is all the extra information that he has in the morning of day 2: a serious weather forecaster, such as the UK Met Office, will have plenty of other information arriving from weather stations around the globe (and even from outer space). The prediction protocol is open in that the forecaster is flooded with information coming from outside the protocol. This is a common situation; to quote [Goldstein \[1983, Sect. 4\]](#), “in most cases of interest (e.g., the doctor’s examination of the patient) it is unreasonable to suppose that, even in principle, there is a partition of possibilities over which probabilities and conditional probabilities could, in theory, be defined.” We will sometimes use the notation \mathcal{F}_{n-1} (formally this is a σ -algebra) for the information available when making the prediction P_n at time n , albeit in many cases this will be an unmanageable notion that is even difficult to imagine (while the moves in our prediction protocols will be observable).

Remark 2.2. Obviously, we can’t include the data arriving from weather stations around the globe in a realistic prediction protocol, but we can go further and argue that even our picture is unrealistic for a large time horizon N : e.g., the first probability measure $P_1 \in \mathfrak{P}(\mathbf{Y}^N)$ specifies $|\mathbf{Y}|^N - 1$ independent parameters, and this number grows exponentially in N even for $|\mathbf{Y}| = 2$. In [Sect. 5](#) we consider a more realistic setting of forecasting K steps ahead (such as a week ahead for $K = 7$).

Remark 2.3. Another reason why we might want to consider a Bayesian violating Bayesian conditioning when updating his beliefs is that his computational resources might be limited: he might keep processing information already available at the previous steps obtaining new values for probabilities of the same events. This is a special case of Hacking’s [\(1967\)](#) observation mentioned earlier.

3 Testing probability forecasts

Long-term prediction is much more complicated than one-step-ahead prediction, and to have a clear understanding of the process we will use two pictures, which we call game-theoretic and measure-theoretic (following [Shafer and Vovk 2019](#)). The game-theoretic picture is based on betting. In this section we discuss the

game-theoretic picture, and in the next one (Sect. 4) we move on to the measure-theoretic picture.

In the game-theoretic picture we add a third player, Sceptic, to the basic forecasting protocol. “Sceptic” is just our name for the better, and betting proceeds according to the following protocol (the intuition behind this protocol will be explained shortly).

Protocol 3.1.

$$\begin{aligned}
& \mathcal{K}_0 := 1 \\
& \text{FOR } n = 1, \dots, N: \\
& \quad \text{Forecaster announces } P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1}) \\
& \quad \text{IF } n > 1: \\
& \quad \quad \mathcal{K}_{n-1} := \mathcal{K}_{n-2} + \sum_{x \in \mathbf{Y}^{N-n+1}} f_{n-1}(y_{n-1}x)P_n(x) \\
& \quad \quad \quad - \sum_{x \in \mathbf{Y}^{N-n+2}} f_{n-1}(x)P_{n-1}(x) \\
& \quad \quad \text{Sceptic announces } f_n : \mathbf{Y}^{N-n+1} \rightarrow \mathbb{R} \\
& \quad \quad \text{Reality announces } y_n \in \mathbf{Y} \\
& \quad \text{IF } n = N: \\
& \quad \quad \mathcal{K}_n := \mathcal{K}_{n-1} + f_n(y_n) - \sum_{y \in \mathbf{Y}} f_n(y)P_n(y).
\end{aligned} \tag{3}$$

Sceptic’s capital \mathcal{K}_n (with an initial capital of 1) is not allowed to become negative (as soon as it does, the play is stopped and Sceptic loses). We regard this protocol (and similar protocols below) as a way of testing Forecaster’s predictions: a large \mathcal{K}_N means lack of agreement of his predictions with reality.

The interpretation of Protocol 3.1 is that at each step n Forecaster announces the price $P_n(x)$ for the uncertain quantity $1_{\{(y_n, \dots, y_N)=x\}}$ for each $x \in \mathbf{Y}^{N-n+1}$. We imagine a ticket that pays $1_{\{(y_n, \dots, y_N)=x\}}$ at the end of the play, and so $P_n(x)$ is Forecaster’s price (at which he prepared to sell and to buy) for this ticket, which we will call *ticket* x . After Forecaster’s move P_n Sceptic buys, for each $x \in \mathbf{Y}^{N-n+1}$, some tickets from Forecaster, and $f_n(x)$ stands for the number of tickets x that he chooses to buy (the number can be positive, negative, or zero). Sceptic pays $\sum_{x \in \mathbf{Y}^{N-n+1}} f_n(x)P_n(x)$ for the transaction. If n is not the last step, at the next step Forecaster announces a new P_n , and Sceptic sells all his tickets at the new prices. Then Sceptic buys a new set of tickets at the new prices, etc. (An important special case is where the new set of tickets coincides with the old set, so effectively no trade happens and Sceptic just keeps the old set.) The change in his capital at step $n < N$ is summarized in (3):

- $\sum_{x \in \mathbf{Y}^{N-n+1}} f_{n-1}(y_{n-1}x)P_n(x)$ is the amount he gains at this step by selling, at the current prices P_n , the tickets that he bought at the previous step; notice that only tickets $y_{n-1}x$ will have non-zero prices;
- $\sum_{x \in \mathbf{Y}^{N-n+2}} f_{n-1}(x)P_{n-1}(x)$ is the amount paid for those tickets at the previous step.

At the last step he just cashes in the winning ticket y_N : see (4).

The next proposition describes the set of admissible moves for Sceptic at step n (see Sect. A.1 for a simple proof). Assume, without loss of generality, $\sum_{x \in \mathbf{Y}^{N-n+1}} f_n(x)P_n(x) = 0$ for all n (there is no loss in generality since additive

constants in f_n cancel out and do not contribute to the change in Sceptic's capital).

Proposition 3.2. *Sceptic's move f_n at step $n = 1, \dots, N$ of Protocol 3.1 guarantees $\mathcal{K}_n \geq 0$ if and only if $f_n \geq -\mathcal{K}_{n-1}$.*

3.1 Comparison with Bayesian conditioning

How does (3) compare with Bayesian conditioning, where no new information outside the protocol arrives and we just define P_n by (2)? In this case we can simplify the protocol by replacing Sceptic's moves f_n with $f'_n : \mathbf{Y} \rightarrow \mathbb{R}$ defined by

$$f'_n(y) := \sum_{x \in \mathbf{Y}^{N-n}} f_n(yx) P_n(x | y), \quad (5)$$

and then (3) becomes

$$\mathcal{K}_{n-1} := \mathcal{K}_{n-2} + f'_{n-1}(y_{n-1}) - \sum_{y \in \mathbf{Y}} f'_{n-1}(y) P_{n-1}(y).$$

Moving this statement to the previous step, we can rewrite Protocol 3.1 as

Protocol 3.3.

```

 $\mathcal{K}_0 := 1$ 
FOR  $n = 1, \dots, N$ :
  IF  $n = 1$ :
    Forecaster announces  $P_1 \in \mathfrak{P}(\mathbf{Y}^N)$ 
  ELSE:
    Forecaster updates  $P_{n-1} \in \mathfrak{P}(\mathbf{Y}^{N-n+2})$  to  $P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1})$ 
    by Bayesian conditioning (2)
  Sceptic announces  $f'_n : \mathbf{Y} \rightarrow \mathbb{R}$ 
  Reality announces  $y_n \in \mathbf{Y}$ 
   $\mathcal{K}_n := \mathcal{K}_{n-1} + f'_n(y_n) - \sum_{y \in \mathbf{Y}} f'_n(y) P_n(y)$ .
```

This is our standard one-step-ahead prediction protocol (cf., e.g., [Shafer and Vovk 2019](#), Protocol 1.1) except that Forecaster announces his forecasting strategy in advance. We can see that forecasting multiple steps ahead does not require any new methods under Bayesian conditioning: testing can proceed one step ahead.

Without any restrictions on Forecaster, we obtain, instead of Protocol 3.3, the following protocol equivalent (in the sense of leading to the same capitals \mathcal{K}_n) to Protocol 3.1, in which we still use the notation f'_n (see (5)).

Protocol 3.4.

```

 $\mathcal{K}_0 := 1$ 
FOR  $n = 1, \dots, N$ :
  Forecaster announces  $P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1})$ 
  IF  $n > 1$ :
```

$$\mathcal{K}_{n-1} := \mathcal{K}_{n-2} + \sum_{x \in \mathbf{Y}^{N-n+1}} f_{n-1}(y_{n-1}x)(P_n(x) - P_{n-1}(x \mid y_{n-1})) \quad (6)$$

Sceptic announces $f_n : \mathbf{Y}^{N-n+1} \rightarrow \mathbb{R}$

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

$$\mathcal{K}_n := \mathcal{K}_{n-1} + f'_n(y_n) - \sum_{y \in \mathbf{Y}} f'_n(y)P_n(y). \quad (7)$$

Protocol 3.4 adds to Protocol 3.3 the possibility to update P_n in a way different from Bayesian conditioning and includes a term that describes betting on the difference between the actual forecast $P_n(x)$ and the Bayesian conditional probabilities $P_{n-1}(x \mid y_{n-1})$ computed from the previous forecast. The equivalence of Protocols 3.1 and 3.4 follows from the equality, for $n < N$, of the addend $f'_n(y_n)$ in line (7) of Protocol 3.4 and the subtrahend $\sum_{x \in \mathbf{Y}^{N-n+1}} f_{n-1}(y_{n-1}x)P_{n-1}(x \mid y_{n-1})$ in line (6) of that protocol at the next step.

3.2 Merging Sceptic's opponents

If we are only interested in strategies for Sceptic (not in strategies for other players, as in Shafer and Vovk 2019, Preface, ideas 3 and 6) we can simplify Protocol 3.1 further by merging Forecaster and Reality. We will refer to the combined player as Forecaster (rather than World, as in Shafer and Vovk 2001, 2019, for agreement with our discussion of Jeffrey's radical probabilism in Sect. D.1 of Appendix D). Let $\mathfrak{P}_n(\mathbf{Y}^N)$ be the set of all probability measures Q on \mathbf{Y}^N satisfying $Q(x) = 1$ for some $x \in \mathbf{Y}^{n-1}$.

Protocol 3.5.

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

FOR $n = 1, \dots, N, N+1$:

Forecaster announces $Q_n \in \mathfrak{P}_n(\mathbf{Y}^N)$

IF $n > 1$:

$$\mathcal{K}_{n-1} := \mathcal{K}_{n-2} + \sum_{x \in \mathbf{Y}^N} F_{n-1}(x)(Q_n(x) - Q_{n-1}(x)) \quad (8)$$

Sceptic announces $F_n : \mathbf{Y}^N \rightarrow \mathbb{R}$.

To embed Protocol 3.1 into Protocol 3.5, we should take as Q_n the extension of P_n to \mathbf{Y}^N , namely

$$Q_n(x) := \begin{cases} P_n(x \setminus (y_1 \dots y_{n-1})) & \text{if } (y_1 \dots y_{n-1}) \subseteq x \\ 0 & \text{otherwise,} \end{cases} \quad (9)$$

and we should take as F_n the extension of f_n to \mathbf{Y}^N , namely

$$F_n(x) := \begin{cases} f_n(x \setminus (y_1 \dots y_{n-1})) & \text{if } (y_1 \dots y_{n-1}) \subseteq x \\ u & \text{otherwise,} \end{cases}$$

where, e.g., $u := 0$ (but in fact the value of u does not matter as it is always multiplied by 0 in the embedded protocol, and we can use different u s for different n and x).

Protocol 3.5 lasts for $N+1$ rather than N steps in order for \mathcal{K}_N to be defined by (8). Sceptic's last move F_{N+1} is never used.

4 Measure-theoretic picture

This section may be skimmed or skipped completely; the remaining sections do not depend on it.

The measure-theoretic picture is stochastic and assumes an overall probability measure used by Forecaster and Reality for generating their moves. In other words, it is just the standard measure-theoretic framework [Kolmogorov, 1933, Doob, 1953] for probability. In this section we will define testing in the measure-theoretic picture and will see that it is equivalent to (albeit more complicated and less natural than) testing in the game-theoretic picture, i.e., the testing procedure described in the previous section. Similar results have been established in game-theoretic probability in the case of one-step-ahead prediction; see, e.g., Shafer and Vovk [2019, Theorem 9.7].

Our check of equivalence will have two sides: the validity of the game-theoretic picture in the measure-theoretic framework, and the validity of a natural measure-theoretic picture in the game-theoretic framework. In our proofs in Appendix A (Sect. A.2) we will use a standard theorem of duality; in general, it can be said that the measure-theoretic and game-theoretic pictures are dual to each other in a certain sense.

What exactly do I mean by equivalence? The idea is to show that we have identical ways of gambling in both pictures. On the measure-theoretic side, we have the standard notion of measure-theoretic martingale (defined later in the section), and we define a *test martingale* as nonnegative martingale with initial value 1. On the game-theoretic side, a *game-theoretic test martingale* is Sceptic's capital \mathcal{K}_n (for all possible n), for a fixed strategy for Sceptic, as function of Forecaster's and Reality's moves provided this function is nonnegative. Roughly, the equivalence means the equivalence of the two notions of test martingale, but the exact statement will become clear in the process of its demonstration (when we reach Proposition 4.1).

4.1 Validity of the game-theoretic picture

Let (Ω, P) be a finite probability space equipped with a filtration \mathcal{F}_n , $n = 0, 1, \dots$. Intuitively, we regard \mathcal{F}_{n-1} as the information available to Forecaster and Sceptic at the beginning of step n in Protocol 3.1. For concreteness, let us assume that all new information (including y_n , which is part of the new information) arrives at the end of step n and none arrives between the steps; therefore, \mathcal{F}_n , $n = 1, 2, \dots$, is the information available at the end of step n .

In the measure-theoretic framework for Protocol 3.1, we assume that y_1, \dots, y_N are realizations of an adapted \mathbf{Y} -valued process Y_1, \dots, Y_N (meaning, as usual, that Y_n is \mathcal{F}_n -measurable, $n = 1, \dots, N$), that each P_n is computed from P as the conditional probability measure for Y_n given \mathcal{F}_{n-1} , and that Sceptic follows a predictable strategy (where “predictable” has the technical meaning that $f_n(x)$ is \mathcal{F}_{n-1} -measurable for each x). Sceptic's capital

\mathcal{K}_n is then an adapted process, and we have

$$P_n(x) = P(\{Y_n \dots Y_N = x\} \mid \mathcal{F}_{n-1}) \quad \text{a.s.,} \quad x \in \mathbf{Y}^{N-n+1}. \quad (10)$$

Now the increment (3) in Sceptic's capital is

$$\begin{aligned} \mathcal{K}_{n-1} - \mathcal{K}_{n-2} &= \sum_{x \in \mathbf{Y}^{N-n+1}} f_{n-1}(Y_{n-1}x) P(\{(Y_n, \dots, Y_N) = x\} \mid \mathcal{F}_{n-1}) \\ &\quad - \sum_{x \in \mathbf{Y}^{N-n+2}} f_{n-1}(x) P(\{(Y_{n-1}, \dots, Y_N) = x\} \mid \mathcal{F}_{n-2}) \\ &= \mathbb{E}_P(f_{n-1}(Y_{n-1}, \dots, Y_N) \mid \mathcal{F}_{n-1}) - \mathbb{E}_P(f_{n-1}(Y_{n-1}, \dots, Y_N) \mid \mathcal{F}_{n-2}), \end{aligned}$$

and so we have

$$\mathbb{E}_P(\mathcal{K}_{n-1} - \mathcal{K}_{n-2} \mid \mathcal{F}_{n-2}) = 0 \quad \text{a.s.} \quad (11)$$

The exceptional (for $n = N$) increment (4) is

$$\mathcal{K}_n - \mathcal{K}_{n-1} = f_n(Y_n) - \mathbb{E}_P(f_n(Y_n) \mid \mathcal{F}_{n-1}),$$

which gives the analogue

$$\mathbb{E}_P(\mathcal{K}_n - \mathcal{K}_{n-1} \mid \mathcal{F}_{n-1}) = 0 \quad \text{a.s.}$$

of (11). Therefore, $\mathcal{K}_0, \mathcal{K}_1, \dots, \mathcal{K}_N$ is a measure-theoretic martingale w.r. to the filtration (\mathcal{F}_n) : $\mathbb{E}_P(\mathcal{K}_n \mid \mathcal{F}_{n-1}) = \mathcal{K}_{n-1}$, $n = 1, \dots, N$.

4.2 Validity of the measure-theoretic picture

A *non-terminal situation* in Protocol 3.1 (and also in Protocol 2.1) is a tuple (P_1, y_1, \dots, P_n) for some $n \in \{1, \dots, N\}$, where $y_i \in \mathbf{Y}$ and $P_i \in \mathfrak{P}(\mathbf{Y}^{N-i+1})$ for all i . Informally, this is a situation in which Sceptic makes a move. A *terminal situation* is a tuple $(P_1, y_1, \dots, P_N, y_N)$, where again $y_i \in \mathbf{Y}$ and $P_i \in \mathfrak{P}(\mathbf{Y}^{N-i+1})$ for all i . Non-terminal situations and terminal situations are referred to collectively as *situations*. A strategy for Sceptic can be defined as a function mapping the non-terminal situations to an allowed move, namely mapping a situation (P_1, y_1, \dots, P_n) to $f : \mathbf{Y}^{N-n+1} \rightarrow \mathbb{R}$ in the case of Protocol 3.1. For a fixed strategy for Sceptic his capital becomes a real-valued function of a situation; as mentioned earlier, we refer to such functions as *game-theoretic test martingales* provided they are nonnegative.

A *game-theoretic process* is a Borel measurable real-valued function of a situation. A nonnegative game-theoretic process S is a *visible measure-theoretic test martingale* if, for any finite probability space (Ω, P) equipped with a filtration $(\mathcal{F}_n)_{n=0}^N$ and any adapted sequence of random variables Y_1, \dots, Y_N ,

$$\begin{aligned} S_n &:= S(P_1, Y_1, \dots, P_{n+1}), \quad n = 0, \dots, N-1, \\ S_N &:= S(P_1, Y_1, \dots, P_N, Y_N) \end{aligned} \quad (12)$$

is a test martingale in the usual sense of $S_0 = 1$ and

$$\mathbb{E}_P(S_n \mid \mathcal{F}_{n-1}) = S_{n-1}, \quad n = 1, \dots, N, \quad (13)$$

where the P_i in (12) are defined by (10), which becomes

$$P_i(x) := P(\{Y_i \dots Y_N = x\} \mid \mathcal{F}_{i-1}), \quad x \in \mathbf{Y}^{N-i+1},$$

in our current notation. The adjective “visible” refers to the martingale (S_n) depending only on the players’ moves in Protocol 2.1 (and not depending on the hidden aspects of the realized sample point $\omega \in \Omega$).

The following statement of agreement between the game-theoretic and measure-theoretic pictures will be proved in Sect. A.2.

Proposition 4.1. *A game-theoretic process is a game-theoretic test martingale if and only if it is a visible measure-theoretic test martingale.*

Proposition 4.1, however, has a weakness. Let us say that a game-theoretic process is a *game-theoretic test supermartingale* if it can be obtained as Sceptic’s capital while he is allowed to discard part of his capital at each step (but is still not allowed to go into debt). For example, in the case of Protocol 3.1 this corresponds to replacing (3) and (4) by Sceptic’s moves allowing him to choose \mathcal{K}_{n-1} and \mathcal{K}_n , respectively, as any nonnegative number not exceeding the corresponding right-hand side. And a game-theoretic process is a *visible measure-theoretic test supermartingale* if it is defined in the same way as a visible measure-theoretic test martingale except that the “=” in (13) is replaced by “ \leq ”. The notion of a game-theoretic test supermartingale is obviously redundant, in the sense of every game-theoretic test supermartingale being dominated by a game-theoretic test martingale. But the requirement (13) holding for any finite probability space might appear restrictive, and so it is less obvious that measure-theoretic test supermartingales are redundant in this sense. Therefore, in Sect. A.2 we will start from proving the following modification of Proposition 4.1.

Theorem 4.2. *A game-theoretic process is a game-theoretic test supermartingale if and only if it is a visible measure-theoretic test supermartingale.*

This theorem implies that every visible measure-theoretic test supermartingale is dominated by a visible measure-theoretic test martingale. We will also check this directly in Sect. A.2.

Remark 4.3. This section illustrates the “hidden variable” account of belief change (Adams, 1975, Chap. 4, note 14, Diaconis and Zabell, 1982, Theorem 2.1, Skyrms, 1992, Sect. 1), according to which belief update is Bayesian conditioning in a bigger belief space.

5 Predicting K steps ahead

For a large time horizon N , the protocols considered in the previous sections are unrealistic in that Forecaster is asked to produce probability measures on huge

sets such as \mathbf{Y}^N . Starting from this section, we will assume that all predictions made by Forecaster are only for the next $K < N$ observations, with $K \geq 1$, and we will sometimes refer to K as the *prediction horizon*. We are typically interested in the case $K \ll N$.

We can still use the Bayesian prediction protocol (Protocol 2.1), but now Sceptic is not allowed to bet more than K steps ahead. In terms of Protocol 3.1, the function $f_n : \mathbf{Y}^{N-n+1} \rightarrow \mathbb{R}$ depends on its argument (y_n, \dots, y_N) only via its first K elements y_n, \dots, y_{n+K-1} (let us assume for the moment that $n + K - 1 \leq N$). Writing $f_n(y_n, \dots, y_{n+K-1})$ instead of $f_n(y_n, \dots, y_N)$, we obtain the following modification of Protocol 3.1.

Protocol 5.1.

```

 $\mathcal{K}_0 := 1$ 
FOR  $n = 1, \dots, N$ :
  Forecaster announces  $P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1})$ 
  IF  $n > 1$ :
     $\mathcal{K}_{n-1} := \mathcal{K}_{n-2} + \sum_{x \in \mathbf{Y}^{(K-1) \wedge (N-n+1)}} f_{n-1}(y_{n-1}x)P_n(x)$ 
     $\quad - \sum_{x \in \mathbf{Y}^{K \wedge (N-n+2)}} f_{n-1}(x)P_{n-1}(x)$ 
  Sceptic announces  $f_n : \mathbf{Y}^{K \wedge (N-n+1)} \rightarrow \mathbb{R}$ 
  Reality announces  $y_n \in \mathbf{Y}$ 
  IF  $n = N$ :
     $\mathcal{K}_n := \mathcal{K}_{n-1} + f_n(y_n) - \sum_{y \in \mathbf{Y}} f_n(y)P_n(y)$ 

```

Of course, we obtain an equivalent protocol if we replace $P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1})$ by $P_n \in \mathfrak{P}(\mathbf{Y}^{K \wedge (N-n+1)})$ in the third line, and this replacement would eliminate an irrelevant part of P_n . Alternatively, we obtain an equivalent protocol if we require $P_n \in \mathfrak{P}(\mathbf{Y}^K)$.

Remark 5.2. In the example of weather forecasting one week ahead (cf. Remark 2.2), the predictions $P_n \in \mathfrak{P}(\mathbf{Y}^K)$ are quite different from the predictions produced by a typical weather app. Weather apps produce marginal probabilities of rain whereas the probabilities in $P_n \in \mathfrak{P}(\mathbf{Y}^K)$ are joint. Testing marginal probabilities is easier than the kind of testing exemplified by Protocol 5.1. See Vovk [2023] for details.

6 Bayesian decision making

Why do we need long-term forecasts? One reason is that they facilitate nearly optimal decisions.

6.1 An optimality result for the Bayes decision strategy

Consider the following decision-making protocol.

Protocol 6.1.

```

FOR  $n = 1, \dots, N$ :
  Reality announces  $\lambda_n : \mathbf{D} \times \mathbf{Y}^{N-n+1} \rightarrow [0, 1]$ 

```

Decision Maker announces $d_n \in \mathbf{D}$
Reality announces the actual observation $y_n \in \mathbf{Y}$.

At each step n Decision Maker is asked to choose a decision d_n from a finite set \mathbf{D} of permitted decisions. Before that, Reality announces a loss function λ_n determining Decision Maker's loss

$$\lambda_n(d_n, y_n \dots y_N) \in [0, 1]$$

at this step. In applications the loss functions are usually given in advance, but we include them in the protocol in order to weaken the conditions of our mathematical result (Theorem 6.5 below). The loss functions are assumed bounded and scaled to the interval $[0, 1]$. The total loss can be computed only after the last step and equals

$$\text{Loss}_N := \sum_{n=1}^N \lambda_n(d_n, y_n \dots y_N) \in [0, N]. \quad (14)$$

Of course, Loss_N is a function of Reality's and Decision Maker's moves, but we will leave the arguments of Loss_N implicit.

A strategy for Decision Maker in Protocol 6.1 is a function giving a decision d_n at each step n as function of Reality's previous moves y_1, \dots, y_{n-1} and $\lambda_1, \dots, \lambda_n$. It would be ideal to have a strategy A for Decision Maker that is provably either better than any other strategy B or approximately equally good, but this is clearly impossible; we need a qualification of the type "with high probability", and our decision making protocol is too poor to express it.

As a first step towards the goal of designing an optimal (in some sense) strategy for Decision Maker, we add a new player, Forecaster, to Protocol 6.1. The following protocol is a combination of Protocols 6.1 and 2.1.

Protocol 6.2.

FOR $n = 1, \dots, N$:
Reality announces $\lambda_n : \mathbf{D} \times \mathbf{Y}^{N-n+1} \rightarrow [0, 1]$
Forecaster announces $P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1})$
Decision Maker announces $d_n \in \mathbf{D}$
Reality announces the actual observation $y_n \in \mathbf{Y}$.

Protocol 6.2 allows us to design a plausible strategy (*Bayes strategy*, or *Bayes optimal strategy*) for Decision Maker (where d_n is now allowed to depend, additionally, on Forecaster's previous moves P_1, \dots, P_n):

$$d_n \in \arg \min_{d \in \mathbf{D}} \sum_{x \in \mathbf{Y}^{N-n+1}} \lambda_n(d, x) P_n(x). \quad (15)$$

However, we cannot prove anything about this strategy as we do not know anything about connections between the forecasts P_n and the actual observations y_n . Therefore, we add Sceptic to our protocol, as in Protocol 3.1.

Protocol 6.3.

$\mathcal{K}_0 := 1$
 FOR $n = 1, \dots, N$:
 Reality announces $\lambda_n : \mathbf{D} \times \mathbf{Y}^{N-n+1} \rightarrow [0, 1]$
 Forecaster announces $P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1})$
 IF $n > 1$:
 $\mathcal{K}_{n-1} := \mathcal{K}_{n-2} + \sum_{x \in \mathbf{Y}^{N-n+1}} f_{n-1}(y_{n-1}x)P_n(x)$
 $- \sum_{x \in \mathbf{Y}^{N-n+2}} f_{n-1}(x)P_{n-1}(x)$
 Decision Maker announces $d_n \in \mathbf{D}$
 Sceptic announces $f_n : \mathbf{Y}^{N-n+1} \rightarrow \mathbb{R}$
 Reality announces $y_n \in \mathbf{Y}$
 IF $n = N$:
 $\mathcal{K}_n := \mathcal{K}_{n-1} + f_n(y_n) - \sum_{y \in \mathbf{Y}} f_n(y)P_n(y)$.

In order to prove a law of large numbers for decision making showing that the Bayes strategy is indeed optimal in some sense, we need the following combination of Protocols 6.3 and 5.1 that only involves prediction K steps ahead. (We will see in Sect. 6.2 that such a law of large numbers inevitably fails for Protocol 6.3.)

Protocol 6.4.

$\mathcal{K}_0 := 1$
 FOR $n = 1, \dots, N$:
 Reality announces $\lambda_n : \mathbf{D} \times \mathbf{Y}^K \rightarrow [0, 1]$
 Forecaster announces $P_n \in \mathfrak{P}(\mathbf{Y}^K)$
 IF $n > 1$:
 $\mathcal{K}_{n-1} := \mathcal{K}_{n-2} + \sum_{x \in \mathbf{Y}^{(K-1) \wedge (N-n+1)}} f_{n-1}(y_{n-1}x)P_n(x)$
 $- \sum_{x \in \mathbf{Y}^{K \wedge (N-n+2)}} f_{n-1}(x)P_{n-1}(x)$ (16)
 Decision Maker announces $d_n \in \mathbf{D}$
 Sceptic announces $f_n : \mathbf{Y}^{K \wedge (N-n+1)} \rightarrow \mathbb{R}$
 Reality announces $y_n \in \mathbf{Y}$
 IF $n = N$:
 $\mathcal{K}_n := \mathcal{K}_{n-1} + f_n(y_n) - \sum_{y \in \mathbf{Y}} f_n(y)P_n(y)$.

We will continue to use the notation Loss_N introduced in (14), which is now modified to

$$\text{Loss}_N := \sum_{n=1}^{N-K+1} \lambda_n(d_n, y_n \dots y_{n+K-1}), \quad (17)$$

but we will also be interested in Decision Maker's loss $\text{Loss}_N(A)$ computed by replacing his actual decisions by the decisions prescribed by a decision strategy A :

$$\text{Loss}_N(A) := \sum_{n=1}^{N-K+1} \lambda_n(d_n^A, y_n \dots y_{n+K-1}),$$

where

$$d_n^A := A(\lambda_1, P_1, y_1, \lambda_2, P_2, \dots, y_{n-1}, \lambda_n, P_n), \quad n = 1, \dots, N - K + 1;$$

we are only interested in strategies that are functions of the previous moves by Reality and Forecaster. Let us adapt the Bayes strategy (15) to Protocol 6.4:

$$d_n := d_n^B \in \arg \min_{d \in \mathbf{D}} \sum_{x \in \mathbf{Y}^K} \lambda_n(d, x) P_n(x), \quad (18)$$

with d_n^B chosen as the first element of the arg min in a fixed linear order on \mathbf{D} if there are ties among d ; $\text{Loss}_N(B)$ is its loss over the first N steps.

If E is a property of Reality's, Forecaster's, and Decision Maker's moves in Protocol 6.4, we define the *upper game-theoretic probability* of E as the infimum of $\alpha > 0$ such that Sceptic has a strategy that guarantees $\mathcal{K}_n \geq 0$ for all n and that ensures $\alpha \mathcal{K}_n \geq 1$ whenever E happens. The following optimality result will be proved in Appendix A (Sect. A.4).

Theorem 6.5. *Let $\epsilon \in (0, 0.3)$. The Bayes strategy B for Decision Maker in Protocol 6.4 guarantees*

$$\bar{\mathbb{P}} \left(\text{Loss}_N(B) - \text{Loss}_N \geq 2\sqrt{KN \ln \frac{1}{\epsilon}} \right) \leq \epsilon. \quad (19)$$

An alternative statement of Theorem 6.5 not using the notion of game-theoretic probability is that there exists a strategy for Sceptic that achieves either

$$\frac{1}{N} \text{Loss}_N(B) - \frac{1}{N} \text{Loss}_N < 2\sqrt{\frac{K}{N} \ln \frac{1}{\epsilon}} \quad (20)$$

or $\mathcal{K}_N \geq 1/\epsilon$. For a small ϵ and large N (as compared with $K \ln \frac{1}{\epsilon}$), this strategy demonstrates that B performs better than or similarly to the actual moves d_n unless Forecaster is discredited. This is a version of the law of large numbers that works only when $K \ll N$.

Remark 6.6. Notice that the strong law of large numbers for a fixed K (and with $N \rightarrow \infty$, as usual) is trivial: we can apply the standard one-step-ahead strong law of large numbers to each K th observation (starting from observation 1, starting from observation 2, ..., and finally starting from observation K). Theorem 6.5 is less trivial, but interestingly, it is based on the same idea. The argument used in the arXiv version 1 of this paper is different but leads to a weaker result (Theorem 7.5 in that version).

Theorem 6.5 shows that, for any other strategy A for Decision Maker, we have

$$\bar{\mathbb{P}} \left(\text{Loss}_N(B) - \text{Loss}_N(A) \geq 2\sqrt{KN \ln \frac{1}{\epsilon}} \right) \leq \epsilon; \quad (21)$$

we, however, prefer the stronger statement (19) allowing Forecaster to choose his moves on the fly. We can rewrite (21) as

$$\bar{\mathbb{P}} \left(\frac{1}{N} \text{Loss}_N(B) - \frac{1}{N} \text{Loss}_N(A) \geq \delta \right) \leq \exp \left(-\frac{\delta^2 N}{4K} \right)$$

for any $\delta \geq 2.2\sqrt{K/N}$. The restriction $\delta \geq 2.2\sqrt{K/N}$ is coming from the condition $\epsilon < 0.3$ in Theorem 6.5; without this restriction, we can still claim that

$$\mathbb{P}\left(\frac{1}{N}\text{Loss}_N(B) - \frac{1}{N}\text{Loss}_N(A) \geq \delta\right) \leq 5\frac{K}{\delta^2 N} \exp\left(-\frac{\delta^2 N}{4K}\right) \quad (22)$$

(for a proof, see the end of Sect. A.4).

Remark 6.7. In Theorem 6.5 we compare Decision Maker's actual loss Loss_N with the loss she would have suffered following the Bayes strategy B , defined by (18). Our interpretation of this theorem depends on the assumption that Reality's and Forecaster's moves are not affected by Decision Maker's moves.

6.2 Predicting $K < N$ steps ahead is essential for our statement of optimality

Theorem 6.5 is about predicting K steps ahead. How important is this restriction? Let us check that it may not be true that

$$\frac{1}{N}(\text{Loss}_N(B) - \text{Loss}_N) < \delta \quad (23)$$

with high probability in Protocol 6.3 for $\delta \ll 1$ if we use the definition of the cumulative loss given in (14) (there is little difference between (14) and (17) for $K \ll N$, but for $K = N$ the latter leads to vacuous statements for $\text{Loss}_N(B) - \text{Loss}_N$); as before, B stands for the Bayes optimal strategy. The intuition behind this demonstration is that at each step Decision Maker is asked to predict the last observation y_N , and this creates heavy dependence between losses at different steps that ruins the law of large numbers.

Set $\mathbf{D} := \mathbf{Y} := \{0, 1\}$, and suppose (in the spirit of measure-theoretic probability) that Forecaster computes his moves from a probability measure $P \in \mathfrak{P}(\{0, 1\}^N)$ governing Reality. The loss functions output by Reality are

$$\lambda_n(d_n, y_n \dots y_N) := \begin{cases} 0 & \text{if } d_n = y_N \\ 1 & \text{otherwise,} \end{cases} \quad (24)$$

and the true probability measure P is such that $P(\{y_N = 1\}) = 0.4$ (so that $y_N = 0$ is slightly likelier than $y_N = 1$).

The Bayes optimal strategy B given by (15) is $d_n^B := 0$. Let us compare it with the complementary strategy $A := 1 - B$ (or simply $A := 1$). We have

$$\frac{1}{N}(\text{Loss}_N(B) - \text{Loss}_N(A)) = \begin{cases} 1 & \text{with probability 0.4} \\ -1 & \text{with probability 0.6,} \end{cases} \quad (25)$$

and so the inequality (23) is grossly violated with a significant probability.

Applying the idea leading to (25) on a smaller scale (to each K th step instead of the last step), we obtain the following lower bound for Protocol 6.4.

Proposition 6.8. *For all N and $K < N/5$,*

$$\mathbb{P} \left(\text{Loss}_N(B) - \text{Loss}_N \geq \sqrt{KN} \right) \geq \epsilon, \quad (26)$$

where B is the Bayes optimal strategy and ϵ is a universal positive constant.

The lower bound \sqrt{KN} in (26) matches the upper bound in (19) (Theorem 6.5) as far as K and N are concerned. (The result in (19) is best interpreted as an upper bound, despite the inequality “ \geq ”; this can be seen from its restatement in the form (20).)

Proposition 6.8 only concerns the optimality of the upper bound in (19) in K and N , but the next proposition shows that it is also close to being optimal in ϵ . In this proposition we use a slightly different definition of Loss_N : now, unlike in (17), we sum the losses of all decisions, including those of d_{N-K+2}, \dots, d_N (they will be defined in a very natural way).

Proposition 6.9. *Suppose that N and K are such that $\sqrt{N/K}$ is an even integer. Then the Bayes optimal strategy B satisfies, for any $\epsilon > 0$ such that $\sqrt{\ln \frac{1}{\epsilon}}$ is integer,*

$$\mathbb{P} \left(\text{Loss}_N(B) - \text{Loss}_N \geq \sqrt{KN \ln \frac{1}{\epsilon}} \right) \geq \epsilon^4/15 \quad (27)$$

provided

$$\sqrt{KN \ln \frac{1}{\epsilon}} \leq N/4. \quad (28)$$

The condition (28) is mild in this context; without it, the bound (27) appears useless. The substitution $\epsilon := \epsilon^4/15$ in Proposition 6.9 gives the following corollary, which shows that the upper bound in (19) is optimal if we ignore additive and multiplicative constants in the “regret term”

$$2\sqrt{KN \ln \frac{1}{\epsilon}}.$$

Corollary 6.10. *Under the conditions of Proposition 6.9,*

$$\mathbb{P} \left(\text{Loss}_N(B) - \text{Loss}_N \geq \frac{1}{2} \sqrt{KN \ln \frac{1}{15\epsilon}} \right) \geq \epsilon \quad (29)$$

provided the term $\sqrt{\dots}$ does not exceed $N/2$.

For proofs of Propositions 6.8 and 6.9, see Sects A.5 and A.6, respectively.

7 Conclusion

This paper has scratched the surface of the diachronic picture of realistic Bayesian forecasting not based on Bayesian conditioning. We discussed ways of testing such forecasts based on betting and their applications to Bayesian decision making.

Obvious directions of further research include, e.g., considering an infinite time horizon and more general observation spaces \mathbf{Y} . Another direction is to generalize our basic forecasting protocol: instead of assuming that the forecaster observes a new outcome y_n at each step, we could consider cases where beliefs are revised (perhaps because new information arrives from outside the protocol) without new outcomes becoming known; for a first step in this direction, see Appendix D.

Acknowledgments

My research has been partially supported by Mitie. Many thanks to Philip Dawid, Ilia Nourtdinov, and participants in the International Seminar on Selective Inference for useful discussions and advice on literature. Reviewers' detailed comments have helped to improve the presentation and are gratefully appreciated.

References

- Ernest W. Adams. *The Logic of Conditionals: An Application of Probability to Deductive Logic*. Springer, Dordrecht, 1975.
- Gunnar Andersson. The problem of the empirical basis in critical rationalism. In Jeremy Shearmur and Geoffrey Stokes, editors, *The Cambridge Companion to Popper*, chapter 5, pages 125–142. Cambridge University Press, New York, 2016.
- Kenneth J. Arrow, Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O. Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D. Nelson, George R. Neumann, Marco Ottaviani, Thomas C. Schelling, Robert J. Shiller, Vernon L. Smith, Erik Snowberg, Cass R. Sunstein, Paul C. Tetlock, Philip E. Tetlock, Hal R. Varian, Justin Wolfers, and Eric Zitzewitz. The promise of prediction markets. *Science*, 320:877–878, 2008.
- Thomas Bayes. An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, 53:370–418, 1763.
- José M. Bernardo and Adrian F. M. Smith. *Bayesian Theory*. Wiley, Chichester, 2000. This is a corrected reissue of a 1994 book.
- Thomas Cover. Universal portfolios. *Mathematical Finance*, 1:1–29, 1991.

- A. Philip Dawid. The well-calibrated Bayesian (with discussion). *Journal of the American Statistical Association*, 77:605–613, 1982.
- A. Philip Dawid. Probability forecasting. In Samuel Kotz, N. Balakrishnan, Campbell B. Read, Brani Vidakovic, and Norman L. Johnson, editors, *Encyclopedia of Statistical Sciences*, volume 10, pages 6445–6452. Wiley, Hoboken, NJ, second edition, 2006.
- A. Philip Dawid and Vladimir Vovk. Prequential probability: Principles and properties. *Bernoulli*, 5:125–162, 1999.
- Bruno de Finetti. La prévision, ses lois logiques, ses sources subjectives. *Annales de l’Institut Henri Poincaré*, 7:1–68, 1937. English translation: Foresight: Its logical laws, its subjective sources. In Kyburg and Smokler [1980] (both first and second editions).
- Bruno de Finetti. *Theory of Probability*. Wiley, Chichester, 2017. First published in 1970 as *Teoria delle probabilità*.
- Persi Diaconis and Sandy L. Zabell. Updating subjective probability. *Journal of the American Statistical Association*, 77:822–830, 1982.
- Joseph L. Doob. *Stochastic Processes*. Wiley, New York, 1953.
- Darrell Duffie. *Futures Markets*. Prentice-Hall, Englewood Cliffs, NJ, 1989.
- Paul Embrechts and Giovanni Puccetti. Bounds for functions of dependent risks. *Finance and Stochastics*, 10:341–352, 2006.
- William Feller. *An Introduction to Probability Theory and Its Applications*, volume 1. Wiley, New York, third edition, 1968.
- Michael Goldstein. The prevision of a prevision. *Journal of the American Statistical Association*, 78:817–819, 1983.
- Peter Grünwald, Rianne de Heide, and Wouter M. Koolen. Safe testing (with discussion). *Journal of the Royal Statistical Society B*, 2024. To appear.
- Ian Hacking. Slightly more realistic personal probability. *Philosophy of Science*, 34:311–325, 1967.
- Larry Harris. *Trading and Exchanges: Market Microstructure for Practitioners*. Oxford University Press, Oxford, 2003.
- John C. Hull. *Options, Futures, and Other Derivatives*. Pearson, Hoboken, NJ, eleventh edition, 2021.
- Richard C. Jeffrey. Radical probabilism (prospectus for a user’s manual). *Philosophical Issues*, 2:193–204, 1992.

- Andrei N. Kolmogorov. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. Springer, Berlin, 1933. English translation: *Foundations of the Theory of Probability*. Chelsea, New York, 1950.
- Robert J. Kunsch and Daniel Rudolf. Optimal confidence for Monte Carlo integration of smooth functions. *Advances in Computational Mathematics*, 45:3095–3122, 2019.
- Henry E. Kyburg, Jr and Howard E. Smokler, editors. *Studies in Subjective Probability*. Krieger, Huntington, NY, second edition, 1980. First edition: Wiley, New York, 1964.
- David Lewis. Why conditionalize? In *Papers in Metaphysics and Epistemology*, chapter 23, pages 403–407. Cambridge University Press, Cambridge, 1999.
- Dennis V. Lindley. *Making Decisions*. Wiley, London, second edition, 1985.
- Jiří Matoušek and Bernd Gärtner. *Understanding and Using Linear Programming*. Springer, Berlin, 2007.
- Jiří Matoušek and Jan Vondrák. The probabilistic method. Available (in August 2025) on [the web](#), 2008.
- Ludovico Piccinato. De finetti’s logic of uncertainty and its impact on statistical thinking and practice. In Prem K. Goel and Arnold Zellner, editors, *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti*, chapter 2, pages 13–30. North-Holland, Amsterdam, 1986.
- Karl R. Popper. *The Logic of Scientific Discovery*. Hutchinson, London, 1950. German original: *Logik der Forschung*, Springer, Vienna, 1934.
- Ulrike Schaeде. Forwards and futures in Tokugawa-period Japan: A new perspective on the Dōjima rice market. *Journal of Banking and Finance*, 13: 487–513, 1989.
- Glenn Shafer. Bayes’s two arguments for the rule of conditioning. *Annals of Statistics*, 10:1075–1089, 1982.
- Glenn Shafer. Conditional probability (with discussion). *International Statistical Review*, 53:261–277, 1985.
- Glenn Shafer. The language of betting as a strategy for statistical and scientific communication (with discussion). *Journal of the Royal Statistical Society A*, 184:407–478, 2021.
- Glenn Shafer. The notion of event in probability and causality: Situating myself relative to Bruno de Finetti. *International Journal of Approximate Reasoning*, 141:171–178, 2022.
- Glenn Shafer and Vladimir Vovk. *Probability and Finance: It’s Only a Game!* Wiley, New York, 2001.

- Glenn Shafer and Vladimir Vovk. *Game-Theoretic Foundations for Probability and Finance*. Wiley, Hoboken, NJ, 2019.
- Albert N. Shiryaev. *Probability-1*. Springer, New York, third edition, 2016.
- Albert N. Shiryaev. *Probability-2*. Springer, New York, third edition, 2019.
- Brian Skyrms. Coherence, probability and induction. *Philosophical Issues*, 2: 215–226, 1992.
- Vladimir Vovk. A logic of probability, with application to the foundations of statistics (with discussion). *Journal of the Royal Statistical Society B*, 55: 317–351, 1993.
- Vladimir Vovk. A game of prediction with expert advice. *Journal of Computer and System Sciences*, 56:153–173, 1998.
- Vladimir Vovk. Logic of subjective probability. Technical Report [arXiv:2309.01173 \[cs.AI\]](#), [arXiv.org](#) e-Print archive, September 2023.
- Vladimir Vovk and Glenn Shafer. A conversation with A. Philip Dawid. *Statistical Science*, 40:148–166, 2025. An extended version is published as [arXiv:2312.00632](#).
- Ian Waudby-Smith and Aaditya Ramdas. Estimating means of bounded random variables by betting (with discussion). *Journal of the Royal Statistical Society B*, 86:1–27, 2024.

A Proofs

A.1 Proof of Proposition 3.2

The condition $f_n \geq -\mathcal{K}_{n-1}$ is sufficient for $\mathcal{K}_n \geq 0$ by (3).

Now suppose $f_n(x) < -\mathcal{K}_{n-1}$ for some $x \in \mathbf{Y}^{N-n+1}$. Fix such an x . If Reality sets y_n to the first element of x and Forecaster sets P_{n+1} to the probability measure on \mathbf{Y}^{N-n} concentrated on $x \setminus y_n$, (3) shows that Sceptic’s capital \mathcal{K}_n will become negative. (Formally, P_{n+1} is required to be positive by our definition of $\mathfrak{P}(\cdot)$, but we can make P_{n+1} arbitrarily close to the probability measure concentrated on $x \setminus y_n$.)

A.2 Proof of Theorem 4.2 and Proposition 4.1

We start from Theorem 4.2. Let us consider in detail only the first step in Protocol 3.1, when we move from prediction $P := P_1$ to prediction $Q := P_2$ (the rest will be easy). We regard P as fixed (so that our argument is conditional on P) and use the notation $P(y)$, where $y \in \mathbf{Y}$, and $P(x \mid y)$, where $y \in \mathbf{Y}$ and $x \in \mathbf{Y}^{N-1}$, as usual. We also use $Q_r(x \mid y)$ for the various $Q = P_2$ possible after observing y as the first observation y_1 (and we assume that Q_r are all

different). We will be able to apply the standard duality theorem since r ranges over a finite set; remember that we consider a finite probability space. It will be convenient to refer to S_1 as the *first value* of a game-theoretic or measure-theoretic martingale $(S_n)_{n=0}^N$.

In Sect. 4.1 we saw that every game-theoretic test martingale is a visible measure-theoretic test martingale, and this implies that every game-theoretic test supermartingale is a visible measure-theoretic test supermartingale; therefore, we will be only interested in the opposite direction. Let $S_{y,r}$ be the first value of a visible measure-theoretic test supermartingale $(S_n)_{n=0}^N$; i.e., $S_{y,r}$ is the first value S_1 when we observe $y_1 = y$ and $P_2 = Q_r$. Our goal is to show that $S_{y,r}$ is the first value of a game-theoretic test supermartingale.

The primary (measure-theoretic) linear programming problem involves variables $X_{y,r} \geq 0$ subject to the constraints

$$\sum_r X_{y,r} = 1$$

for all y and

$$\sum_r X_{y,r} Q_r(x | y) = P(x | y) \quad (30)$$

for all y and x . The interpretation is that $X_{y,r}$ is the conditional probability of Q_r after observing y . The relevant optimization problem is

$$\sum_y P(y) \sum_r S_{y,r} X_{y,r} \rightarrow \max. \quad (31)$$

By the choice of S , the max value is at most 1.

The dual (game-theoretic) problem is

$$\sum_{y,x} P(x | y) Y_{y,x} + \sum_y Y_y \rightarrow \min \quad (32)$$

subject to

$$\sum_x Q_r(x | y) Y_{y,x} + Y_y \geq P(y) S_{y,r} \quad (33)$$

for all y, r . (The recipe for stating the dual problem given in [Matoušek and Gärtner 2007](#), Sect. 6.2, is particularly convenient in this context.) The dual variables $Y_{y,x}$ and Y_y are unconstrained. Rewriting (32) and (33) as

$$\begin{aligned} \sum_{y,x} P(x | y) (Y_{y,x} + Y_y) &\rightarrow \min \\ \sum_x Q_r(x | y) (Y_{y,x} + Y_y) &\geq P(y) S_{y,r}, \end{aligned}$$

respectively, we can see that the optimization problem (32)–(33) is equivalent to

$$\sum_{y,x} P(x | y) Y_{y,x} \rightarrow \min \quad \text{subject to} \quad \sum_x Q_r(x | y) Y_{y,x} \geq P(y) S_{y,r}. \quad (34)$$

Replacing the variables $Y_{y,x}$ with new variables $Z_{y,x}$ defined by $Y_{y,x} = P(y)Z_{y,x}$, we rewrite the optimization problem (34) as

$$\sum_{y,x} P(yx)Z_{y,x} \rightarrow \min \quad \text{subject to} \quad \sum_x Q_r(x | y)Z_{y,x} \geq S_{y,r}, \quad (35)$$

with the same value, at most 1. Any solution to the optimization problem (35) achieves our goal: setting $f_1(yx) := Z_{y,x}$, our portfolio of tickets will have the total final price at least S_1 while their total initial price will be at most 1. To complete the proof of Theorem 4.2, we need to apply the same argument conditionally on the first n observations y_1, \dots, y_n for $n = 1, \dots, N-1$.

Before proving Proposition 4.1 let us make a short detour and check that every visible measure-theoretic test supermartingale is dominated by a visible measure-theoretic test martingale. First we make $S = (S_{y,r})$ admissible replacing each $S_{y,r}$ by the left-hand side of the constraint in (35). The expression being maximized in (31) becomes

$$\begin{aligned} \sum_y P(y) \sum_r X_{y,r} S_{y,r} &= \sum_y P(y) \sum_r X_{y,r} \sum_x Q_r(x | y) Z_{y,x} \\ &= \sum_y P(y) \sum_x Z_{y,x} P(x | y) = \sum_{y,x} P(yx) Z_{y,x}, \end{aligned}$$

where the second equality uses (30). The last expression is very natural, and does not depend at all on the primary variables $X_{y,r}$, which shows that $S_{y,r}$ is the first value of a visible measure-theoretic test martingale except that its initial value can be below 1 (in which case it can be scaled up to make its initial value equal to 1).

Finally, if $(S_{y,r})$ is the first value of a visible measure-theoretic test martingale, it will coincide with the first value of a game-theoretic test supermartingale, which will be the first value of a game-theoretic test martingale (otherwise we could increase this game-theoretic test supermartingale to obtain a visible measure-theoretic martingale whose first value would strictly dominate the first value of the original visible measure-theoretic test martingale, which is impossible). Repeating this argument for y_2, \dots, y_N completes the proof of Proposition 4.1.

A.3 Game-theoretic probability

In the proof of Theorem 6.5 in Sect. A.4 we will need some basic definitions and results in game-theoretic probability given in this subsection; see Shafer and Vovk [2019] for further information. We will let \mathbb{E}_n denote the game-theoretic expectation (to be defined momentarily) at the point in Protocol 6.4 right after Decision Maker announcing her move d_n (let us call this point the *checkpoint*). In our current context \mathbb{E}_n can be defined as follows. If $f = f(y_n \dots y_{(n+K-1) \wedge N})$ is a function of the K consecutive moves by Reality starting from y_n (and ending

with y_N if $n + K - 1 \geq N$),

$$\mathbb{E}_n f := \sum_{x \in \mathbf{Y}^{K \wedge (N-n+1)}} f(x) P_n(x).$$

More generally, if f depends on other future moves (by Reality and other players), $\mathbb{E}_n f$ is the initial capital (if it exists) starting from which Sceptic can attain exactly the final capital of f at the end of step N . If f also depends on the moves preceding the step n checkpoint, $\mathbb{E}_n f$ is found separately for each set of these preceding moves.

The *game-theoretic sample space* Ω consists of all possible sequences of moves

$$\omega := (\lambda_1, P_1, d_1, y_1, \dots, \lambda_N, P_N, d_N, y_N)$$

by non-Sceptic players in Protocol 6.4. A *nonnegative variable* X is a function $X : \Omega \rightarrow [0, \infty)$. The *upper expectation* of X is defined as

$$\overline{\mathbb{E}}(X) := \inf \{ \alpha > 0 \mid \exists \text{ strategy for Sceptic } \forall \omega \in \Omega : \alpha \mathcal{K}_N(\omega) \geq X(\omega) \},$$

where ω are the non-Sceptic player's moves and \mathcal{K}_N is regarded as function of ω . In words, $\overline{\mathbb{E}}(X)$ is the smallest (in the sense of inf) initial capital that Sceptic can turn into $X(\omega)$ or more. An *event* is a set $E \subseteq \Omega$. The *upper probability* $\overline{\mathbb{P}}(E)$ of an event E is defined to be $\overline{\mathbb{E}}(1_E)$.

Lemma A.1. *For any bounded nonnegative variable X ,*

$$\overline{\mathbb{E}}(X) \leq \int_0^\infty \overline{\mathbb{P}}(X \geq u) du. \quad (36)$$

Proof. Set $f(u) := \overline{\mathbb{P}}(X \geq u)$; then $f : [0, \infty) \rightarrow [0, 1]$ is a decreasing function. Replace the ∞ in (36) by C for some upper bound C for X . For each $k = 0, \dots, \lceil C/\epsilon \rceil$, fix a strategy for Sceptic that turns $f(k\epsilon) + \epsilon$ into $1_{\{X \geq k\epsilon\}}$ or more. Multiplying this strategy and its initial capital by ϵ and then summing over the k , we obtain a strategy that turns

$$\sum_{k=0}^{\lceil C/\epsilon \rceil} \epsilon(f(k\epsilon) + \epsilon) \quad (37)$$

into at least

$$\sum_{k=0}^{\lceil C/\epsilon \rceil} \epsilon 1_{\{X(\omega) \geq k\epsilon\}} \geq X(\omega).$$

It remains to notice that (37) tends to $\int_0^C f(u) du$ as $\epsilon \rightarrow 0$. \square

A.4 Proof of Theorem 6.5

This subsection uses the definitions and results from game-theoretic probability given in Sect. A.3. We will also need the following lemma, which is widely used in robust risk aggregation (and our use of this lemma will mimic its uses in robust risk aggregation).

Lemma A.2. For any $C > 0$, any $\alpha \in (0, C/K)$, and any $x_1, \dots, x_K \in \mathbb{R}$,

$$\sum_{k=1}^K g(x_k) \geq 1_{\{\sum_{k=1}^K x_k \geq C\}}, \quad (38)$$

where g is the continuous function

$$g(x) := \begin{cases} 0 & \text{if } x < \frac{C}{K} - \alpha \\ \frac{x - (C/K - \alpha)}{K\alpha} & \text{if } \frac{C}{K} - \alpha \leq x \leq \frac{C}{K} + (K-1)\alpha \\ 1 & \text{if } x > \frac{C}{K} + (K-1)\alpha. \end{cases} \quad (39)$$

Proof. We argue indirectly. Suppose there is a set of numbers x_1, \dots, x_K for which (38) holds with “ $<$ ” in place of “ \geq ”, and let us fix such a set. If $x_i < C/K - \alpha$ and $x_j > C/K + (K-1)\alpha$, we can replace x_i by $x_i + t$ and x_j by $x_j - t$, where $t > 0$ is the smallest number such that $x_i + t = C/K - \alpha$ or $x_j - t = C/K + (K-1)\alpha$; therefore, we can assume, without loss of generality, that there is no such pair (i, j) . In this case, $x_k \leq C/K + (K-1)\alpha$ for all k , but perhaps $x_j < C/K - \alpha$ for some j . It remains to apply Jensen’s inequality to the convex (and increasing) function $g|_{(-\infty, C/K + (K-1)\alpha]}$: as the average of x_k is at least C/K , the average of $g(x_k)$ is at least $g(C/K) = 1/K$. \square

See the proof of Theorem 4.2 in [Embreehts and Puccetti \[2006\]](#) for another proof of Lemma A.2.

Set $Q := \lfloor N/K \rfloor$. Let us first assume that $N = QK + K - 1$; later we will get rid of this assumption (it will be easy as $N = QK + K - 1$ is, in a sense, the worst case).

To get a handle on the difference $\text{Loss}_N(B) - \text{Loss}_N$ in Protocol 6.4, we first consider its increment

$$\lambda_i(d_i^B, y_i \dots y_{i+K-1}) - \lambda_i(d_i, y_i \dots y_{i+K-1}) \quad (40)$$

on step $i \leq N - K + 1$, where d_i^B is the prediction output by the strategy A defined by (18). By the choice of d_i^B , the difference (40) is a supermartingale difference, meaning that its \mathbb{E}_i expectation is nonpositive. Namely,

$$\begin{aligned} \mathbb{E}_i(\lambda_i(d_i^B, y_i \dots y_{i+K-1}) - \lambda_i(d_i, y_i \dots y_{i+K-1})) \\ = \sum_{x \in \mathbf{Y}^K} (\lambda_i(d_i^B, x) - \lambda_i(d_i, x)) P_i(x) \leq 0. \end{aligned}$$

For each $k \in \{1, \dots, K\}$, we consider the process

$$\begin{aligned} L_n^k = \mathbb{E}_n \sum_{i \in \{k, k+K, \dots, k+(Q-1)K\}} \left(\lambda_i(d_i^B, y_i \dots y_{i+K-1}) - \lambda_i(d_i, y_i \dots y_{i+K-1}) \right. \\ \left. + \sum_{x \in \mathbf{Y}^K} (\lambda_i(d_i, x) - \lambda_i(d_i^B, x)) P_i(x) \right); \quad (41) \end{aligned}$$

including only every K th step in the sum simplifies the analysis and, more importantly, makes the result stronger (cf. Remark 6.6). This process starts from zero, and it is a game-theoretic martingale (namely, $L_n^k = \mathcal{K}_{n-1}$ for some strategy for Sceptic in the modification of Protocol 6.4 replacing $\mathcal{K}_n := 1$ by $\mathcal{K}_n := 0$ and allowing \mathcal{K} to become negative), as the following explicit expression shows:

$$\begin{aligned}
L_n^k := & \sum_{i \in \{k, k+K, \dots, k+(q-1)K\}} \left(\lambda_i(d_i^B, y_i \dots y_{i+K-1}) - \lambda_i(d_i, y_i \dots y_{i+K-1}) \right. \\
& \left. + \sum_{x \in \mathbf{Y}^K} (\lambda_i(d_i, x) - \lambda_i(d_i^B, x)) P_i(x) \right) \\
& + \sum_{x \in \mathbf{Y}^{K-j}} (\lambda_{k+qK}(d_{k+qK}^B, y_{k+qK} \dots y_{n-1}x) - \lambda_{k+qK}(d_{k+qK}, y_{k+qK} \dots y_{n-1}x)) P_n(x) \\
& + \sum_{x \in \mathbf{Y}^K} (\lambda_{k+qK}(d_{k+qK}, x) - \lambda_{k+qK}(d_{k+qK}^B, x)) P_{k+qK}(x) \quad (42)
\end{aligned}$$

where q and $j \in \{0, \dots, K-1\}$ are the integers from the representation $n = k + qK + j$, and we are only interested in $n \leq QK$. The first sum (i.e., the sum $\sum_{i \in \{k, k+K, \dots, k+(q-1)K\}}$) in (42) includes the terms (40) (for $i \equiv k \pmod{K}$) that are determined by the checkpoint on step n . The rest of the expression in (42) accounts for the term (40) that is partially determined, which corresponds to $i = k + qK$. And we do not have terms corresponding to $i > k + qK$ since at the checkpoint on step n the expectation of the expression in the outer parentheses in (41) is still 0 for such i .

To check that (42) is indeed a game-theoretic martingale, it suffices to notice that

$$L_n^k - L_{n-1}^k = \sum_{x \in \mathbf{Y}^{K-j}} f_{n-1}(y_{n-1}x) P_n(x) - \sum_{x \in \mathbf{Y}^{K-j+1}} f_{n-1}(x) P_{n-1}(x),$$

where

$$f_{n-1}(x) := \lambda_{k+qK}(d_{k+qK}^B, y_{k+qK} \dots y_{n-2}x) - \lambda_{k+qK}(d_{k+qK}, y_{k+qK} \dots y_{n-2}x),$$

has the same form as the capital increment in (16). This assumes that n is not one of the borderline values $k + qK$, which case should be considered separately.

If we only consider the values of the game-theoretic martingale (42) at steps $k + qK$, $q = 0, 1, \dots, Q$,

$$\begin{aligned}
L_{k+qK}^k := & \sum_{i \in \{k, k+K, \dots, k+(q-1)K\}} \left(\lambda_i(d_i^B, y_i \dots y_{i+K-1}) - \lambda_i(d_i, y_i \dots y_{i+K-1}) \right. \\
& \left. + \sum_{x \in \mathbf{Y}^K} (\lambda_i(d_i, x) - \lambda_i(d_i^B, x)) P_i(x) \right), \quad q = 0, 1, \dots, Q, \quad (43)
\end{aligned}$$

its increments will be bounded by 2 in absolute value, and we can apply the game-theoretic Hoeffding inequality [Shafer and Vovk, 2019, Corollary 3.8 for Protocol 3.5] to it. However, a tighter inequality is obtained when we apply the one-sided version of the game-theoretic Hoeffding inequality [Shafer and Vovk, 2019, Corollary 3.8 for Protocol 3.7] to the process (43) with the sum over $x \in \mathbf{Y}^K$ removed. This process is a game-theoretic supermartingale whose increments are bounded by 1 in absolute value, and the one-sided Hoeffding inequality gives

$$\bar{\mathbb{P}}(X_k \geq U) \leq \exp\left(-\frac{U^2}{2Q}\right) \leq \exp\left(-U^2 \frac{K}{2N}\right), \quad (44)$$

where $U \geq 0$ and

$$X_k := \sum_{i \in \{k, k+K, \dots, k+(Q-1)K\}} (\lambda_i(d_i^B, y_i \dots y_{i+K-1}) - \lambda_i(d_i, y_i \dots y_{i+K-1}));$$

we assume that the game-theoretic supermartingale is constant after $k+(Q-1)K$ (the last i in the range of summation in (43)).

Applying (44) and Lemmas A.1 and A.2 (see below for details) gives, for any $C > 0$,

$$\bar{\mathbb{P}}(X_1 + \dots + X_K \geq C) \leq \sum_{k=1}^K \bar{\mathbb{E}}(g(X_k)) \leq \sum_{k=1}^K \int_0^\infty \bar{\mathbb{P}}(g(X_k) \geq u) du \quad (45)$$

$$\leq \sum_{k=1}^K \int_0^\infty \bar{\mathbb{P}}\left(X_k \geq \gamma \frac{C}{K} + (1-\gamma)Cu\right) du \quad (46)$$

$$\leq \sum_{k=1}^K \int_0^\infty \exp\left(-\left(\gamma \frac{C}{K} + (1-\gamma)Cu\right)^2 \frac{K}{2N}\right) du \quad (47)$$

$$= \frac{\sqrt{KN}}{(1-\gamma)C} \int_{\frac{\gamma C}{\sqrt{KN}}}^\infty \exp(-v^2/2) dv \quad (48)$$

$$= \frac{\sqrt{KN}}{(1-\gamma)C} \sqrt{2\pi} \bar{\Phi}\left(\frac{\gamma C}{\sqrt{KN}}\right) < \frac{KN}{\gamma(1-\gamma)C^2} \exp\left(-\frac{\gamma^2 C^2}{2KN}\right). \quad (49)$$

The first and second inequalities in (45) follow from Lemmas A.2 and A.1, respectively. The inequality (46) follows from the definition of g in (39) with $\alpha := (1-\gamma)C/K$. Indeed, we can assume, without loss of generality, $u > 0$, and then $g(X) \geq u$ implies

$$\frac{X - (C/K - \alpha)}{K\alpha} \geq u,$$

which is equivalent to

$$X \geq \gamma \frac{C}{K} + (1-\gamma)Cu.$$

The inequality (47) follows from Hoeffding's inequality (44). The equality (48) follows by the substitution

$$v := \frac{\gamma C}{\sqrt{KN}} + (1 - \gamma)C\sqrt{\frac{K}{N}}u.$$

The equality in (49) introduces the notation $\bar{\Phi} := 1 - \Phi$ for the survival function of the standard Gaussian distribution. And the last inequality in the chain follows by applying the standard upper bound [Feller, 1968, Lemma VII.1.2] on $\bar{\Phi}$.

We can rewrite the inequality between the extreme terms in the chain (45)–(49) as

$$\bar{\mathbb{P}}(\text{Loss}_N(B) - \text{Loss}_N \geq C) \leq \frac{KN}{\gamma(1 - \gamma)C^2} \exp\left(-\frac{\gamma^2 C^2}{2KN}\right). \quad (50)$$

Comparing this with (19), we can see that we need to solve the inequality

$$\frac{KN}{\gamma(1 - \gamma)C^2} \exp\left(-\frac{\gamma^2 C^2}{2KN}\right) \leq \epsilon. \quad (51)$$

Ignoring the part before the exp and replacing “ \leq ” by “ $=$ ”, we obtain the solution

$$C = \frac{\sqrt{2KN \ln \frac{1}{\epsilon}}}{\gamma},$$

which motivates the substitution

$$C := \frac{\sqrt{2KN \ln \frac{1}{\epsilon} x}}{\gamma} \quad (52)$$

in (51). After this substitution, (51) simplifies to

$$\epsilon^{x-1} \leq 2 \frac{1 - \gamma}{\gamma} x \ln \frac{1}{\epsilon}. \quad (53)$$

Setting $x := 2\gamma^2$ in (52) gives an expression that matches the corresponding expression in (19).

The condition $x > 1$ (required for (53) to hold as $\epsilon \rightarrow 0$) narrows down the range of γ from $\gamma \in (0, 1)$ to $\gamma \in (2^{-1/2}, 1)$. Setting, e.g., $\gamma := 0.8$ ensures that (53) holds for all $\epsilon \in (0, 0.32)$.

It remains to consider the case $N < QK + K - 1$. If the final value of L_{k+qK}^k (corresponding to $q = Q$) is undefined (because $k + qK + K - 1 > N$), we set it equal to its previous value (for $q = Q - 1$).

This completes the proof of Theorem 6.5. Inequality (22) follows from (50) with $\gamma := 1/\sqrt{2}$.

A.5 Proof of Proposition 6.8

Similarly to (24), let us set $\mathbf{D} := \mathbf{Y} := \{0, 1\}$ and

$$\lambda_n(d_n, y_n \dots y_{n+K-1}) := \begin{cases} 1 & \text{if } d_n \neq y_{\lceil n/K \rceil K} \\ 0 & \text{otherwise.} \end{cases} \quad (54)$$

We are only interested in $n \leq N - K + 1$ (see (17)), which implies $n + K - 1 \leq N$ and $\lceil n/K \rceil K \leq N$; therefore, (54) is well-defined. Now the true probability measure P is such that $y_n = 1$ with probability $1/2$ independently for different n (and now we will rely on our tie-breaking convention). As in Sect. 6.2, the players comply with P . Let A be the decision strategy that always outputs 1; notice that B always outputs 0 (assuming that the linear order on \mathbf{D} is $0 < 1$). It suffices to prove (26) (and later (27)) with Loss_N replaced by $\text{Loss}_N(A)$, and this is what we will do.

The N steps are now split into $\lceil N/K \rceil$ blocks of K steps (except, possibly, the last block), $n \in \{1, \dots, K\}$, $n \in \{K + 1, \dots, 2K\}$, etc. Within each block, A suffers the same loss at each step, and B suffers the same loss at each step. By the central limit theorem, the probability is at least ϵ (a universal positive constant) that B performs worse than A in at least $\sqrt{N/K} + 1$ more blocks than vice versa. In such cases

$$\text{Loss}_N(B) - \text{Loss}_N(A) \geq K\sqrt{N/K} = \sqrt{KN}.$$

This gives (26) with P in place of $\bar{\mathbb{P}}$. By Ville's inequality, we can replace the probability measure P by the upper game-theoretic probability $\bar{\mathbb{P}}$.

A.6 Proof of Proposition 6.9

We will obtain Proposition 6.9 by applying a lower bound for large deviations in the form of Matoušek and Vondrák [2008, Proposition 7.3.2] to the argument of the previous subsection. As mentioned before the statement of the proposition, now we define Loss_N by summing the losses over all steps $n = 1, \dots, N$, as in (14). The loss at each step n is still given by the right-hand side of (54). This is the derivation (see below for some explanations):

$$\begin{aligned} & \bar{\mathbb{P}} \left(\text{Loss}_N(B) - \text{Loss}_N(A) \geq \sqrt{KN \ln \frac{1}{\epsilon}} \right) \\ &= \bar{\mathbb{P}} \left(\frac{1}{K} \text{Loss}_N(B) \geq \frac{N}{2K} + \frac{1}{2} \sqrt{\frac{N}{K} \ln \frac{1}{\epsilon}} \right) \end{aligned} \quad (55)$$

$$= \bar{\mathbb{P}} \left(X \geq \frac{m}{2} + t \right) \geq \frac{1}{15} \exp(-16t^2/m) = \epsilon^4/15. \quad (56)$$

The first equality, (55), follows from

$$\text{Loss}_N(A) + \text{Loss}_N(B) = N \quad (57)$$

(which allows us to eliminate $\text{Loss}_N(A)$) and obvious transformations. The first equality in (56) introduces the notation

$$X := \frac{1}{K} \text{Loss}_N(B), \quad m := \frac{N}{K}, \quad t := \frac{1}{2} \sqrt{\frac{N}{K} \ln \frac{1}{\epsilon}},$$

which is close to the notation used in Matoušek and Vondrák [2008, Proposition 7.3.2]. The inequality “ \geq ” in (56) is identical to Matoušek and Vondrák [2008, Proposition 7.3.2].

Matoušek and Vondrák [2008] have two conditions in their Proposition 7.3.2: $t \leq n/8$, which becomes (28), and t being an integer, which we strengthen to $\sqrt{N/K}$ being an even integer and $\sqrt{\ln \frac{1}{\epsilon}}$ being an integer.

Remark A.3. Kunsch and Rudolf [2019, Lemma 3] slightly improve the constants in Matoušek and Vondrák [2008, Proposition 7.3.2], and using their result we can improve the bound $\epsilon^4/15$ in (56) to $\epsilon^3/5$. This allows us to rewrite (29) in the form

$$\mathbb{P} \left(\text{Loss}_N(B) - \text{Loss}_N(A) \geq \sqrt{\frac{1}{3} K N \ln \frac{1}{5\epsilon}} \right) \geq \epsilon. \quad (58)$$

Remark A.4. Let us check informally what the optimal counterparts of the constants $1/3$ and 5 in (58) would be in the domain of applicability of the central limit theorem. We have for the probability measure P of Sect. A.5:

$$P \left(\text{Loss}_N(B) - \text{Loss}_N(A) \geq \bar{\Phi}^{-1}(\epsilon) \sqrt{KN} \right) \approx \epsilon,$$

where $\bar{\Phi}^{-1}(\epsilon)$ is the upper ϵ -quantile of the standard Gaussian distribution. This follows from the variance of

$$\text{Loss}_N(B) - \text{Loss}_N(A) = 2\text{Loss}_N(B) - N$$

(cf. (57)) being approximately KN . This gives the ideal approximate equality

$$P \left(\text{Loss}_N(B) - \text{Loss}_N(A) \geq \sqrt{2KN \ln \frac{1}{\epsilon}} \right) \approx \epsilon$$

in place of (58). It is interesting that this is exactly what we get from (52) when we make $\gamma \approx 1$ and $x \approx 1$ (it is clear that we can make $\gamma \in (0, 1)$ and $x > 1$ as close to 1 as we want at the price of restricting ϵ to a narrower range $(0, \epsilon^*)$).

B Protocols in terms of ideal futures

Testing by betting in general and game-theoretic probability in particular can be interpreted in terms of trading in a financial market, and in this appendix I will make this interpretation explicit. Now we complement the basic forecasting protocol with an idealized market allowing Sceptic to trade in futures contracts

(these are the most standard financial derivatives; see, e.g., [Hull 2021](#), Chap. 2, and [Duffie 1989](#)). Futures contracts is an old idea (see, e.g., [Schaede 1989](#)) that arose gradually in financial industry, but in our prediction protocols it is a powerful way of reducing prediction multiple steps ahead to one-step-ahead prediction. In this appendix we will only need a highly idealized picture of them (Sect. [B.1](#)), but later (Appendix [C](#)) we will discuss their real-life counterparts.

We will also need another piece of notation: $\mathbf{Y}^{m:n}$ stands for the set of all sequences of elements of \mathbf{Y} of length between m and n inclusive (so that $\mathbf{Y}^{0:n}$ stands for the sequences of elements of \mathbf{Y} of length at most n , and $\mathbf{Y}^{1:n}$ stands for the non-empty sequences of elements of \mathbf{Y} of length at most n).

B.1 Ideal futures contracts

In Sect. [3](#) we extended the forecasting picture of Sect. [2](#) by allowing Sceptic to bet against Forecaster. Betting was described in terms of tickets, which are known as forward contracts in finance. In our diachronic picture, however, we allowed trade in tickets (they were sold and bought), which essentially turned them into what is known as futures contracts in finance. In this appendix we will talk about futures contracts explicitly using very convenient terminology developed in finance. Our terminology, however, will be slightly adapted to our needs (for example, the unit of time will be a step rather than, e.g., a day, and the trader will be called Sceptic).

A futures contract Φ has an *expiration step* m . The contract is settled at the end of step m ; namely, its final price F_m^+ is announced by Reality. In the middle of step $n \in \{1, \dots, m\}$, the current price F_n of Φ is announced by Forecaster, and Sceptic can then take any *position* $f_n \in \mathbb{R}$ in Φ . If $n < m$, Sceptic gains capital $f_n(F_{n+1} - F_n)$ at the next step (which actually means losing capital if $f_n(F_{n+1} - F_n) < 0$). If $n = m$, at the end of the expiration step m (at *maturity*) Sceptic gains $f_m(F_m^+ - F_m)$. These gains keep accumulating as the play proceeds.

Let us say that Sceptic takes a *constant position* f at time $n < m$ if he maintain the same position f through steps n, \dots, m . This leads to gaining capital $f(F_m^+ - F_n)$ at maturity. This mode of using futures contracts emulates forward contracts (which are similar to futures contracts but not exchange-traded).

B.2 General testing protocol

The following extension of Protocol [2.1](#) describes another, seemingly stronger than Protocol [3.1](#), way of testing Forecaster's predictions. (However, we will reduce the extended protocol to Protocol [3.1](#) in Sect. [B.3](#).)

Protocol B.1.

$\mathcal{K}_0 := 1$
Forecaster announces $P_1 \in \mathfrak{P}(\mathbf{Y}^N)$
Sceptic announces $f_1 : \mathbf{Y}^{1:N} \rightarrow \mathbb{R}$

$$\begin{aligned}
& \text{Reality announces } y_1 \in \mathbf{Y} \\
& \mathcal{K}'_1 := \mathcal{K}_0 + f_1(y_1) - \sum_{y \in \mathbf{Y}} f_1(y) P_1(y) \\
& \text{FOR } n = 2, \dots, N: \\
& \quad \text{Forecaster announces } P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1}) \\
& \quad \mathcal{K}_{n-1} := \mathcal{K}'_{n-1} + \sum_{x \in \mathbf{Y}^{1:(N-n+1)}} f_{n-1}(y_{n-1}x) P_n(x) \\
& \quad \quad - \sum_{x \in \mathbf{Y}^{2:(N-n+2)}} f_{n-1}(x) P_{n-1}(x) \\
& \quad \text{Sceptic announces } f_n : \mathbf{Y}^{1:(N-n+1)} \rightarrow \mathbb{R} \\
& \quad \text{Reality announces } y_n \in \mathbf{Y} \\
& \quad \mathcal{K}'_n := \mathcal{K}_{n-1} + f_n(y_n) - \sum_{y \in \mathbf{Y}} f_n(y) P_n(y).
\end{aligned} \tag{59}$$

Protocol [B.1](#) does not define \mathcal{K}_N , and we set $\mathcal{K}_N := \mathcal{K}'_N$. The interpretation of \mathcal{K}_N is the same as for Protocol [3.1](#): a large \mathcal{K}_N evidences lack of agreement of the forecasts with reality, provided Sceptic is not allowed to go into debt.

An advantage of Protocol [B.1](#) over Protocol [3.1](#) is that, even though it is stated for a finite time horizon N , it is easier to modify to make the time horizon infinite, so that $n = 2, 3, \dots$ in the FOR loop. The simpler protocol of Sect. [3](#) uses the finiteness of the time horizon in a more essential way.

The financial interpretation of Protocol [B.1](#) is that we have a market of futures contracts $\Phi(x)$, $x \in \mathbf{Y}^{1:N}$, that pay

$$F_m^+(x) := 1_{\{y_1 \dots y_m = x\}}$$

at the end of step $m := |x|$, as discussed in Sect. [B.1](#). At each step n (but before observing y_n) Forecaster announces the prices for all the futures contracts $\Phi(x)$, $x \in \mathbf{Y}^{1:N}$, in the form of a probability measure $P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1})$; namely, the price of $\Phi(x)$, $x \in \mathbf{Y}^{1:N}$, at step n is

$$F_n(x) := \begin{cases} P_n(x \setminus y_1 \dots y_{n-1}) & \text{if } y_1 \dots y_{n-1} \subset x \\ 0 & \text{if not.} \end{cases} \tag{61}$$

A standard argument shows that such P_n will exist provided the market is coherent; Sceptic can secure a sure gain if F_n do not form a probability measure concentrated on the continuations of $y_1 \dots y_{n-1}$ [[de Finetti, 2017](#), Chap. 3].

At step n Sceptic needs to take positions in all $\Phi(x)$, $y_1 \dots y_{n-1} \subset x \in \mathbf{Y}^{1:N}$. The position in $\Phi(y_1 \dots y_{n-1}x)$ is denoted by $f_n(x)$ in Protocol [B.1](#). (There is no need to take positions in the other $\Phi(x)$ since their prices are 0 and will stay 0.)

After y_n is disclosed by Reality, the increment in Sceptic's capital (due to the matured futures contracts $\Phi(y_1 \dots y_{n-1}y)$) is

$$\begin{aligned}
\mathcal{K}'_n - \mathcal{K}_{n-1} &= \sum_{y \in \mathbf{Y}} f_n(y) (F_n^+(y_1 \dots y_{n-1}y) - F_n(y_1 \dots y_{n-1}y)) \\
&= f_n(y_n) - \sum_{y \in \mathbf{Y}} f_n(y) P_n(y),
\end{aligned}$$

which agrees with [\(60\)](#). And after P_n is disclosed by Forecaster at the next step $n := n + 1$, the increment in Sceptic's capital (due to the remaining futures

contracts) is

$$\begin{aligned}
\mathcal{K}_{n-1} - \mathcal{K}'_{n-1} &= \sum_{x \in \mathbf{Y}^{2:(N-n+2)}} f_{n-1}(x) (F_n(y_1 \dots y_{n-2}x) - F_{n-1}(y_1 \dots y_{n-2}x)) \\
&= \sum_{x \in \mathbf{Y}^{1:(N-n+1)}} f_{n-1}(y_{n-1}x) F_n(y_1 \dots y_{n-1}x) \\
&\quad - \sum_{x \in \mathbf{Y}^{2:(N-n+2)}} f_{n-1}(x) F_{n-1}(y_1 \dots y_{n-2}x) \\
&= \sum_{x \in \mathbf{Y}^{1:(N-n+1)}} f_{n-1}(y_{n-1}x) P_n(x) - \sum_{x \in \mathbf{Y}^{2:(N-n+2)}} f_{n-1}(x) P_{n-1}(x),
\end{aligned}$$

where the last equality follows from (61), and the last expression agrees with (59).

B.3 Simplification

We can rewrite Protocol B.1 in other forms, such as Protocol 3.1, getting rid of some of Sceptic's arbitrary choices. To compare protocols with the same allowed moves for Reality and Forecaster, we can use the notion of the *test martingale space* (TMS), which we define modifying the definition given in Sect. 4 as follows. A strategy for Sceptic still specifies his move as function of Forecaster's and Reality's previous moves, but now we do not impose any measurability conditions on strategies. As before, the corresponding *test martingale* is Sceptic's capital as function of Forecaster's and Reality's moves provided this function is nonnegative. The TMS of a given protocol is the set of all possible test martingales. We regard two protocols to be equivalent if they have the same TMS.

As already mentioned, the general testing protocol, Protocol B.1, was formulated with a view towards an infinite time horizon, where N becomes ∞ . In Sect. 3 we introduced a much simpler protocol using an idea that only works for a finite time horizon.

Proposition B.2. *Protocol B.1 and Protocol 3.1 have identical TMS.*

Proposition B.2 simplifies the market in futures contracts that we need: all the contracts now mature at the end of step N ; we will call such futures contracts *final*. The intuitive reason why the final futures contracts are sufficient is that a general futures contract $\Phi(x)$ is equivalent, to all intents and purposes, to the portfolio consisting of the final futures contracts $\Phi(x')$ for all $x' \supseteq x$.

Proof of Proposition B.2. Consider step $n < N$ of Protocol B.1. Let $O(x, c)$, where $x \in \mathbf{Y}^{1:(N-n)}$ and $c \in \mathbb{R}$, be the operation that adds the constant c to $f_n(x)$ and subtracts the same constant c from all $f_n(xy)$, $y \in \mathbf{Y}$. The key observation used in our simplification of Protocol B.1 is that, for any $x \in \mathbf{Y}^{1:(N-n)}$ and $c \in \mathbb{R}$, $O(x, c)$ does not change the increment in the capital $\mathcal{K}_n - \mathcal{K}_{n-1}$. Let us check this property. If $x \in \mathbf{Y}^{2:(N-n)}$, $O(x, c)$ will not affect

(60) whatsoever, and it will change neither minuend nor subtrahend in (59) at the next step (there is a next step since $n < N$). And if $x \in \mathbf{Y}$, applying the operation $O(x, c)$ does not change the increment in the capital $\mathcal{K}_n - \mathcal{K}_{n-1}$ given by (60) and then by (59) at the next step since

- the changes in the sum in (60) and in the second sum in (59) at the next step will balance each other out, and
- the changes in the term $f_n(y_n)$ in (60) and in the first sum in (59) at the next step will also balance each other out (this is relevant only when $x = y_n$).

Applying $O(x, c)$ repeatedly to the x s in the order of increasing length, we can assume, without loss of generality (i.e., without changing the TMS), that $f_n(x)$ is different from 0 only for $x \in \mathbf{Y}^{N-n+1}$, which implies that:

- we can ignore (60) for all steps n apart from $n = N$, and so (4) is performed only for $n = N$;
- we can ignore the bits “1 : ” and “2 : ” in (59), obtaining (3).

Protocol 3.1 also merges the four lines in Protocol B.1 preceding the FOR loop into the loop. \square

C Mechanics of futures trading

Section B.1 gives an idealized picture of futures trading. The main elements of simplification in it are:

- the interest rate is assumed to be zero;
- the positions and futures prices are assumed to take any real values (although we are only interested in positive prices for futures contracts);
- there is no difference between the selling and buying prices (no bid/ask spread);
- there are no other transaction costs.

In this paper we are only interested in binary futures contracts (where the outcome is 0 or 1). However, the most popular market mechanism, described in this appendix, works for general futures contracts, which are not restricted to the binary case.

A good reference for traditional futures markets is Duffie [1989]. While some of the physical details of trading described in it might be obsolete, the general principles are still applicable. Another good reference is Harris [2003].

By far the most popular platform for prediction markets is the Iowa Electronic Markets (IEM). The IEM was created in 1988 and has always been a small-scale operation; the development of prediction markets has been greatly

hindered by the US anti-gambling regulation [Arrow et al., 2008]. The IEM was created by academics, and its role is mainly educational; in particular, it has a great help system explaining the market microstructure (which I often follow in this section). It received two no-action letters, in 1992 and 1993, from the US Commodity Futures Trading Commission (CFTC) reducing the chance of legal action against it. Its competitors sometimes have better bid/ask spreads, but their positions are less secure; e.g., Intrade (1999–2013) is now defunct and PredictIt (launched in 2014) had their CFTC no-action letter withdrawn in 2022.

A futures contract is a contract that pays a specified amount F_m at a specified future time, called the *expiration time* m (it was the expiration step in the main part of the paper). The amount is uncertain at the time of trading but becomes well-defined at the expiration time, when trading ceases. An example of m and F_m is “6 November 2024” and “Democratic Nominee’s share of the two-party popular vote in the 2024 US Presidential election” in US dollars. This is, essentially, one of the types of futures contracts traded at the IEM in August 2023 (of the “vote share” variety; the other main variety is “winner takes all”). Let us fix m and F_m . At each time the market participants can hold any number of the futures contracts (positive, zero, or negative), which is known as their *positions* in the futures contracts. They can also submit orders to change their positions. The main kinds of orders are *market orders* and *limit orders*. A limit order specifies the number of futures contracts to buy or sell at a given price (known as the *bid price* for orders to buy and the *ask price* for orders to sell); it may also specify the time when the order expires.

At the core of a futures market is the *order book* listing the outstanding limit orders. The prices specified in those orders are

$$B_{n_B} < B_{n_B-1} < \cdots < B_1 < A_1 < A_2 < \cdots < A_{n_A}, \quad (62)$$

where n_B is the number of different bid prices in the currently active limit orders to buy and n_A is the number of different ask prices in the currently active limit orders to sell. The prices in the list (62) are sorted in the ascending order, and the difference $A_1 - B_1$ is known as the *bid/ask spread*. With each price level x is associated the total number $N(x)$ of futures contracts that the market participants with active limit orders wish to trade (to buy if $x = B_n$ for some n and to sell if $x = A_n$ for some n ; $N(x) = 0$ for all other x). The order book consists of the prices (62) and the numbers $N(x)$ of futures contracts offered at each price level x (within each price level x older orders appear before newer orders). It consists of a *bid queue* (the data related to the bid prices) and an *ask queue* (the data related to the ask prices).

A market order is simpler than a limit order and only specifies the number of futures contracts to buy or sell. When a new market order is submitted by a market participant MP, it is matched with the order book immediately and a trade is performed. Namely, if the order is to sell N contracts, the bid queue is traversed from the top (i.e., from B_1) until the required number of orders to buy is found: we find the smallest k such that $N(B_1) + \cdots + N(B_k) \geq N$ (all the

$N(B_1) + \dots + N(B_{n_B})$ contracts requested in the bid queue are bought if there is no such k) and arrange a trade with MP selling all his futures contracts to the market participants with active limit orders with the prices in $\{B_1, \dots, B_k\}$; for the price B_k only the oldest orders are fulfilled (perhaps partially). The procedure for market orders to buy is analogous.

When a new limit order is submitted by a market participant, it is simply added to the order book. We can assume that the limit orders to buy specify prices below A_1 and the limit orders to sell specify prices above B_1 (otherwise, a market order can be submitted). When a limit order in the order book expires, it is, of course, removed from it.

An important element of futures markets is the system of *margins*. Typically market participants have positions in several types of futures contracts (corresponding to different m and F_m) and other securities, and the total values of their portfolios can go up or down. To reduce the chance of the exchange losing money, they are required to maintain margin accounts at specified levels. If a margin account falls below the specified level as result of changing market prices, a *margin call* is issued requiring the account to be replenished.

In the IEM, short (i.e., negative) positions are formally prohibited, which allows it to avoid imposing margin requirements. But it is still easy to emulate short positions (e.g., a short position in the vote share for the Democratic Nominee can be modelled as a long position in the vote share for the Republican Nominee).

A natural question is how a futures market is started; namely how to make the order book non-empty. In the IEM, the market participants are allowed to buy *fixed price bundles* for a given price. For example, such a bundle might contain the vote share for the Democratic Nominee and the vote share for the Republican Nominee, with a fixed price of \$1 (the sum of the two vote shares is 1, and so the final pay-off of the bundle is known to be \$1).

D Radical probabilism

Our testing protocols, such as Protocol 3.1, assume that we learn the observations y_n with full certainty. According to Jeffrey’s doctrine of radical probabilism [Jeffrey, 1992], we do not learn anything for certain; at best, we learn that the n th observation is y_n with a high probability. Radical probabilism can be regarded as extension of Cromwell’s rule, which is an assumption about the synchronic picture at each step n , to the diachronic picture. The uncertainty of observations is a recurring topic in the philosophy of science. See, e.g., Popper’s discussion of “basic statements” in Popper [1950, Chap. 5] (where he also refers to Reininger’s and Neurath’s similar ideas) and Andersson [2016]. In this section we will discuss two modifications of Protocol 3.1 allowing uncertain evidence.

D.1 Additive picture

A straightforward modification of Protocol 3.5 making evidence uncertain is the following one.

Protocol D.1.

$$\begin{aligned}
&\mathcal{K}_0 := 1 \\
&\text{FOR } n = 1, 2, \dots : \\
&\quad \text{Forecaster announces } Q_n \in \mathfrak{P}(\mathbf{Y}^N) \\
&\quad \text{IF } n > 1: \\
&\quad \quad \mathcal{K}_{n-1} := \mathcal{K}_{n-2} + \sum_{x \in \mathbf{Y}^N} F_{n-1}(x)(Q_n(x) - Q_{n-1}(x)) \\
&\quad \text{Sceptic announces } F_n : \mathbf{Y}^N \rightarrow \mathbb{R}.
\end{aligned} \tag{63}$$

Whereas the loops in Protocols 3.1 and 3.5 are over finite ranges of n , in Protocol D.1 the loop is infinite since we do not learn any of y_1, \dots, y_N with certainty. Even though in Protocol D.1 y_n are never disclosed explicitly, they may be disclosed implicitly via Q_n : cf. (9). The capital updating rule (63) is very natural: namely, a possible interpretation of this rule is that Q_{n-1} is the expectation of Q_n (cf. Goldstein 1983, Theorem in Sect. 3).

Notice that Protocol D.1 is only a modification, not generalization, of Protocol 3.5: whereas Q_n in the former protocol is required to be positive (by our definition of $\mathfrak{P}(\cdot)$), it is not positive in the latter protocol (being positive is incompatible with possessing certain evidence).

D.2 Multiplicative picture

Protocol D.1 and all the protocols discussed in the main part of the paper are similar to Protocol 3.1 in that Sceptic's capital is updated by adding various terms. This subsection introduces a multiplicative protocol, in which Sceptic's capital is updated by multiplication. Both multiplicative and additive protocols are ubiquitous in game-theoretic probability (although the difference between them is rarely pointed out). This is the multiplicative version of Protocol D.1:

Protocol D.2.

$$\begin{aligned}
&\mathcal{K}_0 := 1 \\
&\text{FOR } n = 1, 2, \dots : \\
&\quad \text{Forecaster announces } Q_n \in \mathfrak{P}(\mathbf{Y}^N) \\
&\quad \text{IF } n > 1: \\
&\quad \quad \mathcal{K}_{n-1} := \mathcal{K}_{n-2} \left(G_{n-1}(\square) + \sum_{x \in \mathbf{Y}^N} \frac{Q_n(x)}{Q_{n-1}(x)} G_{n-1}(x) \right) \\
&\quad \text{Sceptic announces } G_n : (\mathbf{Y}^N \cup \{\square\}) \rightarrow \mathbb{R} \text{ summing to 1.}
\end{aligned} \tag{64}$$

This protocol uses the notation \square for the empty sequence, and in its last line Sceptic is required to ensure $G_n(\square) + \sum_{x \in \mathbf{Y}^N} G_n(x) = 1$. The interpretation of (64) is that Sceptic invests some of his capital \mathcal{K}_{n-2} into tickets x for various $x \in \mathbf{Y}^N$ and deposits the remainder of \mathcal{K}_{n-2} into a zero-interest bank account \square ; the share of the capital invested or deposited is $G_{n-1}(x)$ and $G_{n-1}(\square)$, respectively. As usual, the capital \mathcal{K}_n is not allowed to become negative.

To see the equivalence of the additive and multiplicative protocols, notice that (64) is equivalent to

$$\begin{aligned}\mathcal{K}_{n-1} - \mathcal{K}_{n-2} &= \left(G_{n-1}(\square) + \sum_{x \in \mathbf{Y}^N} \frac{Q_n(x)}{Q_{n-1}(x)} G_{n-1}(x) - 1 \right) \mathcal{K}_{n-2} \\ &= \left(\sum_{x \in \mathbf{Y}^N} (Q_n(x) - Q_{n-1}(x)) \frac{G_{n-1}(x)}{Q_{n-1}(x)} \right) \mathcal{K}_{n-2}.\end{aligned}$$

This establishes the one-to-one correspondence

$$F_{n-1}(x) = \frac{G_{n-1}(x)}{Q_{n-1}(x)} \mathcal{K}_{n-2} \quad (65)$$

between F_{n-1} in (63) and G_{n-1} in (64). The correspondence (65) assumes that $\mathcal{K}_{n-2} > 0$, and the case $\mathcal{K}_{n-2} = 0$ should be considered separately (Sceptic's capital will stay at 0 once it reaches 0 in either protocol).

We obtain a useful modification of Protocol D.2 replacing $G_n : (\mathbf{Y}^N \cup \{\square\}) \rightarrow \mathbb{R}$ in the last line by $G_n \in \mathfrak{P}(\mathbf{Y}^N)$. Then the multiplicative protocol becomes a special case of Cover's protocol modelling investment into $|\mathbf{Y}|^N$ securities such as stocks (see, e.g., Cover 1991 or Vovk 1998, Example 9). As in Sect. 3, we have a market in securities $\Phi(x)$, $x \in \mathbf{Y}^N$, but they may be never settled. For each security $\Phi(x)$ the protocol gives its price $Q_n(x)$ at time n . The prices are normalized in that $Q_n(x)$ sum to 1 over x ; e.g., $Q_n(x)$ may be the market shares. The capital update rule (64) involves the *price relative* $Q_n(x)/Q_{n-1}(x)$ (as used in Cover 1991). At each step Sceptic decides on the distribution G_n of his current capital \mathcal{K}_{n-1} among the securities $\Phi(x)$. If $G_n \in \mathfrak{P}(\mathbf{Y}^N)$, we do not allow “short selling”, i.e., holding a negative amount of a security, and we require Sceptic to invest all of his capital. In general, allowing any $G_n : \mathbf{Y}^N \rightarrow \mathbb{R}$ we allow both short selling and leaving part (positive or negative) of Sceptic's capital on a zero-interest bank account.

D.3 Radical probabilism and reality

The additive picture and, especially, the multiplicative one shed new light on the protocols in the main part of the paper. The latter cover the case where Q_n , $n = 1, \dots, N$, is concentrated on $[x] \subseteq \mathbf{Y}^N$ (the set of all continuations of x) for some $x \in \mathbf{Y}^{N-n+1}$. The difference between radical probabilism and the standard Bayesian scenario considered in the main paper corresponds to the difference between stocks and futures contracts. Sooner or later, reality settles a futures contract, but stock prices can be forever variable (in our ideal picture).

It would be interesting to establish conditions under which this paper's results can be extended to the more general and simpler protocols of this appendix.

E Internal and external coherence

A fundamental role in Bayesian statistics is played by the requirement of *coherence*: if the Bayesian’s beliefs do not form a probability measure, we can set a “Dutch book” against him, which is a system of bets leading to his sure loss. In this paper we were mainly interested in a stronger property, agreement with reality, in which sure loss is replaced by a substantial gain for an opponent (Sceptic) who follows a strategy that can ever lose only a small amount (1 in our prediction protocols). Both coherence and agreement with reality are defined in terms of betting.

One subtlety of the standard Bayesian treatment of coherence is that coherence only implies finite additivity and not countable additivity (see, e.g., [de Finetti 2017](#), Sect. 18.3, and [Bernardo and Smith 2000](#), Sect. 3.5.2). In the finite case considered in this paper, however, the difference between finite and countable additivity disappears. (Starting from a finite case is standard in probability theory; see, e.g., [Kolmogorov 1933](#), Chap. I, and [Shiryaev 2016](#), Chap. 1.)

Coherence is a property of consistency of the Bayesian’s beliefs, so we could call it internal coherence. The English word “coherence” may cover not only internal coherence but also agreement with reality (a kind of external coherence). For example, one of the earliest abstract uses of “coherence” given in the Oxford English Dictionary is from Abraham Fraunce’s “The lawiers logike” (1588): “Where there is a greater cohærence and affinitie betweene the argument and the thing argued”. While using “coherence” in its meaning of internal coherence is standard in Bayesian statistics, here we will also discuss external coherence. I will argue that Lewis’s diachronic Dutch book argument is a step towards external coherence and that the game-theoretic testing procedure involving Sceptic can also be interpreted as a check of external coherence.

The following proposition will show that neither the standard static notion of coherence nor Lewis’s dynamic requirement of no diachronic Dutch book impose any restrictions on the forecasts and observations at different times. Therefore, for discussing the diachronic aspects of Bayesian forecasting we need a stronger requirement, such as withstanding Sceptic’s tests. And the latter requirement appears to be a natural variation on Lewis’s dynamic coherence.

Proposition E.1. *For any sequence of outcomes $y_1, \dots, y_N \in \mathbf{Y}$ and any sequence of probability measures $P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1})$, $n = 1, \dots, N$, there is a positive finite probability space (Ω, P) with filtration $\mathcal{F}_0, \dots, \mathcal{F}_N$ and an adapted sequence of \mathbf{Y} -valued random elements Y_1, \dots, Y_N such that the event*

$$\forall n \in \{1, \dots, N\} : P_n = P(\cdot \mid \mathcal{F}_{n-1}) \ \& \ Y_n = y_n$$

has a positive probability.

The Bayesian is incoherent in the usual sense of Ramsey and de Finetti if we can set up a system of bets under which he always loses, called a Dutch book. This static notion is not directly applicable in the diachronic setting since we

learn the full sequence P_1, \dots, P_N only at the very end of the forecasting session, when it is too late to bet.

But we can, in principle, apply Lewis's dynamic modification of the usual static notion of coherence. This modification was used by Lewis [1999, p. 406] in his diachronic Dutch book argument; lack of coherence becomes, in Lewis's words, "a risk of loss uncompensated by any chance of gain". Unlike Lewis, however, we do not assume that our prediction protocol is comprehensive; otherwise, we would have to conclude, as Lewis did, that P_n are connected via Bayesian conditioning.

Let us say that the Bayesian (Forecaster in our protocol) is *dynamically incoherent* in a particular play if there is a way of betting against him that never leads to the better's loss but, in this particular play, leads to the better's gain. In other words, it's a gain that is not compensated by a potential loss; we will call it a *gratis gain*. Lewis [1999] (and other people before 1999) used the expression "diachronic Dutch book" to mean the existence of a play for which the Bayesian is dynamically incoherent. While dynamic coherence is external (depends on the actual outcomes y_i), the requirement of no diachronic Dutch book becomes one of internal coherence.

Proposition E.1 implies that Forecaster is never dynamically incoherent provided the bets are fair (the expectation of the better's final capital does not exceed the initial capital) under every possible probability measure P compatible with the forecasts. This is true under any strategy for probability updating (or in the absence of such a strategy). Indeed, if the gain in the capital is positive for play $P_1 y_1 \dots P_N y_N$ and nonnegative for all other plays, the expected gain will be positive under any positive probability measure P compatible with the play $P_1 y_1 \dots P_N y_N$, and the existence of such a P is asserted in Proposition E.1; therefore, the bets cannot be fair.

We can see that there can be no diachronic inconsistency between P_n for different n or between P_n and y_n leading to a gratis gain for the better. Our testing protocol in Sect. 3, however, shows that such inconsistency can lead to an **almost** gratis gain for the better. Namely, when for a particular play \mathcal{K}_N is large, we can regard it as an almost gratis gain for the better (Sceptic). (This assumes that Sceptic's capital is measured in small monetary units, but we can always scale it up if the monetary units are not small.) For example, a gross failure of the law of large numbers would lead to an almost gratis gain.

Proof of Proposition E.1. Let us fix such sequences of probability measures $P_n \in \mathfrak{P}(\mathbf{Y}^{N-n+1})$ and outcomes $y_n \in \mathbf{Y}$. As a first step, define Ω as \mathbf{Y}^N , P as P_1 , $Y_n(\omega)$ as the n th element ω_n of $\omega \in \Omega$, and let the σ -algebra \mathcal{F}_n be generated by Y_1, \dots, Y_n .

Next modify the finite probability space (Ω, P) and filtration (\mathcal{F}_n) as follows (there is no need to modify Y_n). Split each sample point $y_1 \omega_2 \dots \omega_N$ that starts from y_1 into two sample points, $y'_1 \omega_2 \dots \omega_N$ and $y''_1 \omega_2 \dots \omega_N$, and make the sets $\{y'_1\} \times \Omega^{N-1}$ and $\{y''_1\} \times \Omega^{N-1}$ \mathcal{F}_n -measurable for $n \geq 1$. Split the old value $c := P_1(\{y_1\} \times \Omega^{N-1})$ into $P(\{y'_1\} \times \Omega^{N-1}) := \epsilon c$, for a sufficiently small $\epsilon > 0$, and $P(\{y''_1\} \times \Omega^{N-1}) := (1 - \epsilon)c$. Without changing $P(\{\omega_1 \dots \omega_N\})$ for

$\omega_1 \notin \{y'_1, y''_1\}$, set

$$\frac{P(\{y'_1 \omega_2 \dots \omega_N\})}{P(\{y'_1\} \times \Omega^{N-1})} := P_2(\{\omega_2 \dots \omega_N\}), \quad \omega_2, \dots, \omega_N \in \Omega,$$

and define

$$\frac{P(\{y''_1 \omega_2 \dots \omega_N\})}{P(\{y''_1\} \times \Omega^{N-1})}, \quad \omega_2, \dots, \omega_N \in \Omega,$$

in such a way that we have an agreement with P_1 :

$$\forall \omega_2, \dots, \omega_N \in \Omega : \frac{P(\{y'_1 \omega_2 \dots \omega_N, y''_1 \omega_2 \dots \omega_N\})}{P(\{y'_1, y''_1\} \times \Omega^{N-1})} = \frac{P_1(\{y_1 \omega_2 \dots \omega_N\})}{P_1(\{y_1\} \times \Omega^{N-1})},$$

this is possible for a sufficiently small $\epsilon > 0$.

Apply the same procedure to the probability subspace of (Ω, P) consisting of the sample points $y'_1 \omega_2 \dots \omega_N$, thereby splitting y_2 into y'_2 and y''_2 . Continue by splitting y_3, y_4 , etc. \square